



Informatica®
10.5

Administrator Guide

Informatica Administrator Guide

10.5

March 2021

© Copyright Informatica LLC 2005, 2023

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, the Informatica logo, PowerCenter, Power Exchange, and Informatica Cloud 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: 2023-07-18

Table of Contents

Preface	15
Informatica Resources.	15
Informatica Network.	15
Informatica Knowledge Base.	15
Informatica Documentation.	15
Informatica Product Availability Matrices.	16
Informatica Velocity.	16
Informatica Marketplace.	16
Informatica Global Customer Support.	16
Chapter 1: Understanding Domains	17
Understanding Domains Overview.	17
Nodes.	18
Service Manager.	18
Application Services.	20
Analyst Service.	21
Content Management Service.	22
Data Integration Service.	22
Email Service.	23
Metadata Access Service.	23
Metadata Manager Service.	23
Model Repository Service.	23
PowerCenter Integration Service.	24
PowerCenter Repository Service.	24
PowerExchange Listener Service.	24
PowerExchange Logger Service.	24
Resource Manager Service.	25
SAP BW Service.	25
Scheduler Service.	25
Web Services Hub.	25
High Availability.	25
Informatica Data Usage Policy.	26
Configuring Informatica DiscoveryIQ Proxy Details.	26
Disabling Informatica Data Usage.	26
Chapter 2: Managing Your Account	28
Managing Your Account Overview.	28
Password Management.	28
Changing Your Password.	29
Preferences.	29

Informatica Network Credentials.	29
Enter Informatica Network Credentials.	30
Searching Informatica Knowledge Base.	30
Chapter 3: Using Informatica Administrator.	31
Using Informatica Administrator Overview.	31
Log In to Informatica Administrator.	32
Informatica Administrator URL.	33
Troubleshooting the Login to Informatica Administrator.	33
Manage Tab.	33
Manage Tab - Domain View.	34
Details Panel.	36
Resource Usage Indicators.	36
Manage Tab - Services and Nodes View.	38
Navigator Search.	39
Domain.	39
Folders.	40
Application Services.	40
System Services.	44
Nodes.	45
Grids.	46
Licenses.	46
Manage Tab - Connections View.	46
Manage Tab - Schedules View.	47
Monitor Tab.	48
Monitor Tab - Summary Statistics View.	49
Monitor Tab - Execution Statistics View.	50
Views in the Execution Statistics View.	51
Statistics in the Execution Statistics View.	52
Reports on the Execution Statistics View.	53
Summary Statistics View.	55
Detailed Statistics View.	56
Logs Tab.	56
Reports Tab.	57
Security Tab.	57
Using the Search Section.	57
Using the Security Navigator.	57
Groups.	58
Users.	58
Roles.	59
Operating System Profiles.	59
LDAP Configuration.	60
Account Management.	60

Audit Reports.	60
Service States.	61
Process States.	61
Job States.	63
Informatica Administrator Accessibility Overview.	63
Keyboard Shortcuts.	64
Chapter 4: Using the Domain View.	65
About the Domain View.	65
Dependency Graph.	66
Viewing Dependencies for Application Services, Nodes, and Grids.	66
Recycling or Disabling Downstream Services.	67
Command History.	68
History View.	68
Viewing History.	69
Viewing Events.	70
Chapter 5: Domain Management.	72
Domain Management Overview.	72
Alert Management.	73
Configuring SMTP Settings.	73
Subscribing to Alerts.	74
Viewing Alerts.	74
Folder Management.	75
Creating a Folder.	75
Moving Objects to a Folder.	76
Removing a Folder.	76
Domain Security Management.	76
User Security Management.	77
Application Service Management.	77
Enabling and Disabling Services and Service Processes.	78
Viewing Service Processes.	78
Configuring Restart for Service Processes.	79
Removing Application Services.	79
Troubleshooting Application Services.	79
Gateway Configuration.	80
Configuring the Gateway and Worker Nodes.	80
Domain Configuration Management.	81
Backing Up the Domain Configuration.	81
Restoring the Domain Configuration.	82
Migrating the Domain Configuration.	82
Updating the Domain Configuration Database Connection.	84
Rename the Domain.	85

Shutting Down a Domain.	85
Domain Properties.	86
General Properties.	86
Database Properties.	87
Gateway Configuration Properties.	88
Service Level Management.	88
SMTP Configuration.	89
Custom Properties for the Domain.	90
Chapter 6: Nodes.	91
Nodes Overview.	91
Node Types.	92
Gateway Nodes.	92
Worker Nodes.	92
Example Domain with Multiple Nodes.	92
Node Roles.	93
Service Role.	93
Compute Role.	94
Updating the Node Role.	94
Viewing Processes on a Node with the Service Role.	95
Define and Add Nodes.	95
Adding Nodes to the Domain.	95
Configuring Node Properties.	96
Shutting Down and Restarting the Node.	98
Shutting Down a Node from the Administrator Tool.	99
Starting or Stopping a Node on Windows.	99
Starting or Stopping a Node on UNIX.	99
Removing the Node Association.	100
Removing a Node.	100
Chapter 7: High Availability.	101
High Availability Overview.	101
Resilience.	102
Application Client Resilience.	102
Application Service Resilience.	103
Node Resilience.	103
Example Resilience Timeout Configuration.	104
Restart and Failover.	105
Domain Failover.	105
Application Service Restart and Failover.	106
Recovery.	106
Configuration for a Highly Available Domain.	107
Application Service Resilience Configuration.	108

Application Service Failover Configuration.	109
PowerCenter Integration Service Failover and Recovery Configuration.	109
Command Line Program Resilience Configuration.	110
Domain Failover Configuration.	110
Node Restart Configuration.	111
Troubleshooting High Availability.	111
Chapter 8: Connections.	112
Connections Overview.	112
Connection Management.	112
Creating a Connection.	113
Refreshing the Connections List.	113
Viewing a Connection.	114
Configuring Pooling for a Connection.	114
Editing and Testing a Connection.	114
Deleting a Connection.	115
Pass-through Security.	115
Pass-Through Security with Data Object Caching.	116
Adding Pass-Through Security	116
Pooling Properties in Connection Objects.	117
Chapter 9: Connection Properties.	118
Connection Properties Overview.	119
Adabas Connection Properties.	120
Amazon Redshift Connection Properties.	122
Amazon S3 Connection Properties.	124
Blockchain Connection Properties.	126
Cassandra Connection Properties.	128
Confluent Kafka Connection.	129
General Properties.	130
Confluent Kafka Broker Properties.	130
SSL Properties	131
Creating a Confluent Kafka Connection Using infacmd.	131
Databricks Connection Properties.	131
Greenplum Connection Properties.	133
Google Analytics Connection Properties.	135
Google BigQuery Connection Properties.	135
Connection Modes.	137
Google Cloud Spanner Connection Properties.	137
Google Cloud Storage Connection Properties.	138
Google PubSub Connection Properties.	139
Hadoop Connection Properties.	140
Hadoop Cluster Properties.	141

Common Properties.	142
Reject Directory Properties.	144
Blaze Configuration.	144
Spark Configuration.	145
HBase Connection Properties.	146
HDFS Connection Properties.	147
HBase Connection Properties for MapR-DB.	149
Hive Connection Properties.	149
HTTP Connection Properties.	153
IBM DB2 Connection Properties.	155
IBM DB2 for i5/OS Connection Properties.	158
IBM DB2 for z/OS Connection Properties.	161
IMS Connection Properties.	164
JDBC Connection Properties.	166
JDBC V2 Connection Properties.	169
JD Edwards EnterpriseOne Connection Properties.	171
Kafka Connection Properties.	172
General Properties.	173
Kafka Broker Properties.	173
SSL Properties.	174
Creating a Kafka Connection Using infacmd.	175
Kudu Connection Properties	175
LDAP Connection Properties.	176
Microsoft Azure Blob Storage Connection Properties.	177
Microsoft Azure Cosmos DB SQL API Connection Properties.	178
Microsoft Azure Data Lake Storage Gen1 Connection Properties.	179
Microsoft Azure Data Lake Storage Gen2 Connection Properties.	180
Microsoft Azure SQL Data Warehouse Connection Properties.	181
MS SQL Server Connection Properties.	182
Netezza Connection Properties.	186
OData Connection Properties.	187
ODBC Connection Properties.	188
Oracle Connection Properties.	189
Salesforce Connection Properties.	192
Salesforce Marketing Cloud Connection Properties.	193
SAP Connection Properties.	194
Sequential Connection Properties.	197
Snowflake Connection Properties.	199
Teradata Parallel Transporter Connection Properties.	200
Tableau Connection Properties.	202
Tableau V3 Connection Properties.	203
Twitter Streaming Connection Properties.	204

VSAM Connection Properties.	205
Web Services Connection Properties.	208
Identifier Properties in Database Connections.	209
Regular Identifiers.	210
Delimited Identifiers.	210
Identifier Properties.	210
Chapter 10: Schedules.	213
Schedules Overview.	213
Create and Edit Schedules.	213
Creating a Schedule.	214
Editing a Schedule.	216
Pausing and Resuming a Schedule.	216
Removing Jobs from a Schedule.	217
Deleting a Schedule.	217
Chapter 11: Domain Object Export and Import.	218
Domain Object Export and Import Overview.	218
Export Process.	218
Rules and Guidelines for Exporting Domain Objects.	219
View Domain Objects.	219
Viewable Domain Object Names.	220
Import Process.	220
Rules and Guidelines for Importing Domain Objects.	221
Conflict Resolution.	221
Chapter 12: License Management.	222
License Management Overview.	222
License Validation.	223
Licensing Log Events.	223
License Management Tasks.	223
Types of License Keys.	224
Original Keys.	224
Incremental Keys.	224
Creating a License Object.	225
Assigning a License to a Service.	226
Rules and Guidelines for Assigning a License to a Service.	226
Unassigning a License from a Service.	226
Updating a License.	227
Removing a License.	227
License Properties.	228
License Details.	228
Supported Platforms.	229

Repositories.	229
Service Options.	230
Connections.	230
Metadata Exchange Options.	230
Chapter 13: Monitoring.	231
Monitoring Overview.	231
Configuring Monitoring.	232
Step 1. Configure Monitoring Settings.	233
Step 2. Configure Reports and Statistics Views.	234
Optimizing Monitoring Performance.	234
Summary Statistics.	235
Viewing Summary Statistics.	236
Monitor Data Integration Services	237
Properties View for a Data Integration Service.	237
Reports View for a Data Integration Service.	237
Monitor Ad Hoc Jobs.	238
Viewing Logs for an Ad Hoc Job.	241
Canceling an Ad Hoc Job.	241
Viewing Summary Statistics for an Ad Hoc Job.	241
Viewing Detailed Statistics for an Ad Hoc Job.	241
Monitoring Mapping Audits.	242
Monitor Applications.	242
Properties View for an Application.	242
Reports View for an Application.	243
Monitor Deployed Mapping Jobs.	243
Viewing Logs for a Deployed Mapping Job.	243
Reissuing a Deployed Mapping Job.	244
Canceling a Deployed Mapping Job.	244
Viewing Summary Statistics for a Deployed Mapping Job.	244
Viewing Detailed Statistics for a Deployed Mapping Job.	245
Viewing Deployed Mapping Job Statistics with the REST Operations Hub Service.	245
Monitor Logical Data Objects.	246
Properties View for a Logical Data Object.	246
Cache Refresh Runs View for a Logical Data Object.	246
Viewing Logs for Data Object Cache Refresh Runs.	246
Monitor SQL Data Services.	247
Properties View for an SQL Data Service.	247
Connections View for an SQL Data Service.	247
Requests View for an SQL Data Service.	248
Virtual Tables View for an SQL Data Service.	249
Reports View for an SQL Data Service.	249
Monitor Web Services.	250

Properties View for a Web Service.	250
Reports View for a Web Service.	250
Operations View for a REST or SOAP Web Service.	250
Requests View for a Web Service.	251
Monitor Workflows.	251
Workflow Graph	251
View Workflow Objects.	252
Workflow States.	253
Workflow Object States.	254
Mapping Task Work Item States.	255
Canceling or Aborting a Workflow.	256
Workflow Recovery.	256
Workflow Logs.	257
Job Status After Application Service Restart or Failover.	259
Monitoring a Folder of Objects.	259
Viewing the Context of an Object.	260
Configuring the Date and Time Custom Filter.	260
Configuring the Elapsed Time Custom Filter.	260
Configuring the Multi-Select Custom Filter.	260
Chapter 14: Log Management.	262
Log Management Overview.	262
Log Manager Architecture.	263
PowerCenter Session and Workflow Log Events.	263
Data Integration Service Job Log Events.	264
Log Manager Recovery.	264
Troubleshooting the Log Manager.	264
Log Location.	265
System Logs.	265
Log Management Configuration.	266
Purging Log Events.	266
Time Zone.	267
Configuring Log Management Properties.	267
Using the Logs Tab.	267
Viewing Log Events.	268
Configuring Log Columns.	269
Saving Log Events.	270
Exporting Log Events.	270
Viewing Administrator Tool Log Errors.	272
Log Events.	272
Log Event Components.	273
Domain Log Events.	273
Analyst Service Log Events.	274

Data Integration Service Log Events.	274
Listener Service Log Events.	274
Logger Service Log Events.	275
Model Repository Service Log Events.	275
Metadata Manager Service Log Events.	275
PowerCenter Integration Service Log Events.	275
PowerCenter Repository Service Log Events.	276
Resource Manager Service Log Events.	276
SAP BW Service Log Events.	276
Scheduler Service Log Events.	277
Web Services Hub Log Events.	277
User Activity Log Events.	277
Log Aggregator.	278
Aggregating Application Service Logs.	278
Processing Aggregated Application Service Logs.	279
Mapping Task Logs.	279
Chapter 15: Domain Reports.	281
Domain Reports Overview.	281
License Management Report.	281
Licensing.	282
CPU Summary.	282
CPU Detail.	283
Repository Summary.	284
User Summary.	284
User Detail.	284
Hardware Configuration.	285
Node Configuration.	285
Licensed Options.	286
Running the License Management Report.	286
Sending the License Management Report in an Email.	287
Web Services Report.	288
Understanding the Web Services Report.	288
General Properties and Web Services Hub Summary.	289
Web Services Historical Statistics.	290
Web Services Run-time Statistics.	290
Web Service Properties.	291
Web Service Top IP Addresses.	291
Web Service Historical Statistics Table.	291
Running the Web Services Report.	292
Running the Web Services Report for a Secure Web Services Hub.	293

Chapter 16: Node Diagnostics.....	294
Node Diagnostics Overview.	294
Informatica Network Login.	295
Logging In to the Informatica Network.	295
Generating Node Diagnostics.	296
Downloading Node Diagnostics.	297
Uploading Node Diagnostics.	297
Analyzing Node Diagnostics.	298
Identify Bug Fixes.	298
Identify Recommendations.	299
Chapter 17: Understanding Globalization.....	300
Globalization Overview.	300
Unicode.	301
Working with a Unicode PowerCenter Repository.	301
Locales.	302
System Locale.	302
User Locale.	303
Input Locale.	303
Data Movement Modes.	303
Character Data Movement Modes.	303
Changing Data Movement Modes.	304
Code Page Overview.	305
UNIX Code Pages.	306
Windows Code Pages.	306
Choosing a Code Page.	307
Code Page Compatibility.	307
Domain Configuration Database Code Page.	308
Administrator Tool Code Page.	309
PowerCenter Client Code Page.	309
PowerCenter Integration Service Process Code Page.	309
PowerCenter Repository Code Page.	310
Metadata Manager Repository Code Page.	310
PowerCenter Source Code Page.	310
PowerCenter Target Code Page.	311
Command Line Program Code Pages.	311
Code Page Compatibility Summary.	312
Code Page Validation.	314
Relaxed Code Page Validation.	315
Configuring the PowerCenter Integration Service.	316
Selecting Compatible Source and Target Code Pages.	316
Troubleshooting for Code Page Relaxation.	316

PowerCenter Code Page Conversion.	317
Choosing Characters for PowerCenter Repository Metadata.	317
Case Study: Processing ISO 8859-1 Data.	318
Configuring the ISO 8859-1 Environment.	318
Case Study: Processing Unicode UTF-8 Data.	320
Configuring the UTF-8 Environment.	320
Appendix A: Code Pages.	323
Supported Code Pages for Application Services.	323
Supported Code Pages for Sources and Targets.	325
Appendix B: Custom Roles.	335
Analyst Service Custom Role.	335
Metadata Manager Service Custom Roles.	336
Operator Custom Role.	337
PowerCenter Repository Service Custom Roles.	338
Test Data Manager Custom Roles.	339
Appendix C: Informatica Platform Connectivity.	343
Informatica Platform Connectivity Overview.	343
Domain Connectivity.	344
Model Repository Connectivity.	345
PowerCenter Connectivity.	346
Repository Service Connectivity.	348
Integration Service Connectivity.	348
PowerCenter Client Connectivity.	349
Metadata Manager Service Connectivity.	350
Native Connectivity.	351
ODBC Connectivity.	351
JDBC Connectivity.	352
Appendix D: Configure the Web Browser.	353
Configure the Web Browser.	353
Index.	354

Preface

Use the *Informatica® Administrator Guide* to learn how to log into the Administrator tool and understand the user interface. Read on how to configure, manage, and monitor the Informatica domain. Learn about domain architecture and its components, including nodes, services, high availability, connections, and monitoring.

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>.

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

Understanding Domains

This chapter includes the following topics:

- [Understanding Domains Overview, 17](#)
- [Nodes, 18](#)
- [Service Manager, 18](#)
- [Application Services, 20](#)
- [High Availability, 25](#)
- [Informatica Data Usage Policy, 26](#)

Understanding Domains Overview

Informatica has a service-oriented architecture that provides the ability to scale services and share resources across multiple machines. High availability functionality helps minimize service downtime due to unexpected failures or scheduled maintenance in the Informatica environment.

The Informatica domain is the fundamental administrative unit in Informatica. The domain supports the administration of the distributed services. A domain is a collection of nodes and services that you can group in folders based on administration ownership.

A node is the logical representation of a machine in a domain. One node in the domain acts as a gateway to receive service requests from clients and route them to the appropriate service and node. Services and processes run on nodes in a domain. The availability of a service or process on a node depends on how you configure the service and the node.

Services for the domain include the Service Manager and a set of application services:

- **Service Manager.** A service that runs on each node in the domain to manage all domain operations. The Service Manager performs domain functions such as authentication, authorization, and logging. The Service Manager also starts the application services configured to run on the node.
- **Application Services.** Services that represent server-based functionality, such as the Model Repository Service and the Data Integration Service. The application services that run on a node depend on the way you configure the services.

The Service Manager and application services control security. The Service Manager manages users and groups that can log in to application clients and authenticates the users who log in to the application clients. The Service Manager and application services authorize user requests from application clients.

Informatica Administrator (the Administrator tool), consolidates the administrative tasks for domain objects such as services, nodes, licenses, and grids. You manage the domain and the security of the domain through the Administrator tool.

If you have the high availability option, you can scale services and eliminate single points of failure for services. Services can continue running despite temporary network or hardware failures.

Nodes

A node is a logical representation of a machine in a domain. During installation, you add the installation machine to the domain as a node. You can add multiple nodes to a domain.

Each node in the domain runs the Service Manager that manages domain functions on that node. The Service Manager also supports the application services that run on the node. The domain functions that the node performs and the services that the node runs depend on the following node configurations:

Node type

The node type determines whether the node can serve as a gateway or worker node and determines the domain functions that the node performs. One gateway node serves as the master gateway node for the domain. The master gateway node receives service requests from clients and routes them to the appropriate service and node. A worker node is any node not configured to serve as a gateway. The first time that you install the Informatica services, you create a gateway node and the Informatica domain. When you install the Informatica services on other machines, you create additional gateway nodes or worker nodes that you join to the domain.

Node role

The node role defines the purpose of the node. A node with the service role can run application services. A node with the compute role can perform computations requested by remote application services. A node with both roles can run application services and locally perform computations for those services. By default, each gateway and worker node has both the service and compute roles enabled. If a node is assigned to a Data Integration Service grid, you might want to update the node role. Enable only the service role to dedicate the node to running the Data Integration Service process. Enable only the compute role to dedicate the node to running Data Integration Service mappings.

You can subscribe to alerts to receive notification about node events such as node failure or a master gateway election. You can also generate and upload node diagnostics to the Configuration Support Manager and review information such as available EBFs and Informatica recommendations.

Service Manager

The Service Manager is a service that manages all domain operations. It runs within Informatica services. It runs as a service on Windows and as a daemon on UNIX. When you start Informatica services, you start the Service Manager.

The Service Manager runs on each node in the domain. If the Service Manager is not running, the node is not available.

The Service Manager runs on all nodes in the domain to support the domain and application services:

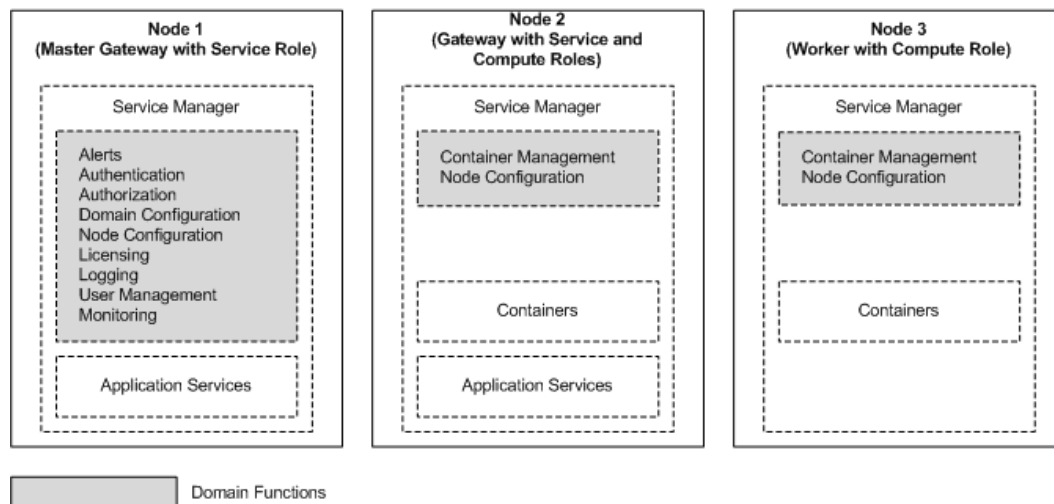
- **Domain support.** The Service Manager performs functions on each node to support the domain. The functions that the Service Manager performs on a node depend on the type and role of the node. For example, the Service Manager running on the master gateway node performs all domain functions on that node. The Service Manager running on any other type of node performs limited domain functions on that node.
- **Application service support.** When a node has the service role, the Service Manager starts application services configured to run on that node. It starts and stops services and service processes based on requests from clients. It also directs service requests to application services. The Service Manager uses TCP/IP to communicate with the application services.

The following table describes the domain functions that the Service Manager performs:

Function	Description
Alert Management	The Service Manager sends alerts to subscribed users. You subscribe to alerts to receive notification for node failure and master gateway election on the domain, and for service process failover for services on the domain. When you subscribe to alerts, you receive notification emails. Alert management occurs on the master gateway node.
Authentication	The Service Manager authenticates users who log in to application clients. Authentication occurs on the master gateway node.
Authorization	The Service Manager authorizes user requests for domain objects based on the privileges, roles, and permissions assigned to the user. Requests for domain objects can come from the Administrator tool. Domain authorization occurs on the master gateway node.
Container Management	When a node has the compute role, the Service Manager manages the containers on the node. A container is an allocation of memory and CPU resources. An application service uses the container to remotely perform computations on the node. For example, a Data Integration Service that runs on a grid can remotely run a mapping within a container on a node with the compute role. Container management occurs on any node with the compute role.
Domain Configuration	The Service Manager manages the domain configuration metadata. Domain configuration occurs on the master gateway node.
Node Configuration	The Service Manager manages node configuration metadata in the domain. Node configuration occurs on all nodes in the domain.
Licensing	The Service Manager registers license information and verifies license information when you run application services. Licensing occurs on the master gateway node.
Logging	The Service Manager provides accumulated log events from each service in the domain. To perform the logging function, the Service Manager runs a Log Manager and a Log Agent. The Log Manager runs on the master gateway node. The Log Agent runs on all nodes where PowerCenter® Integration Service sessions and workflows run and where Data Integration Service jobs run.

Function	Description
User Management	The Service Manager manages the native and LDAP users and groups that can log in to application clients. It also manages the creation of roles and the assignment of roles and privileges to native and LDAP users and groups. User management occurs on the master gateway node.
Monitoring	The Service Manager persists, updates, retrieves, and publishes run-time statistics for integration objects in the monitoring Model repository. The Service Manager stores the monitoring configuration details in the domain configuration repository. Monitoring occurs on the master gateway node.

The following image shows where the Service Manager performs domain functions:



Application Services

Application services represent server-based functionality.

Application services include services that you create and system services that are created for you when you create the domain. A system service can have a single instance in the domain.

Application services include the following services:

- Analyst Service
- Content Management Service
- Data Integration Service
- Metadata Access Service
- Metadata Manager Service
- Model Repository Service
- PowerCenter Integration Service
- PowerCenter Repository Service

- PowerExchange® Listener Service
- PowerExchange Logger Service
- Resource Manager Service
- SAP BW Service
- Scheduler Service
- Test Data Manager Service
- Web Services Hub

When you configure an application service, you designate a node to run the service process. When a service process runs, the Service Manager assigns a port number from the range of port numbers assigned to the node.

The service process is the runtime representation of a service running on a node. The service type determines how many service processes can run at a time. For example, the PowerCenter Integration Service can run multiple service processes at a time when you run it on a grid.

If you have the high availability option, you can run a service on multiple nodes. Designate the primary node to run the service. All other nodes are back-up nodes for the service. If the primary node is not available, the service runs on a back-up node. You can subscribe to alerts to receive notification in the event of a service process failover.

If you do not have the high availability option, configure a service to run on one node. If you assign the service to multiple nodes, then the service will not start.

Analyst Service

The Analyst Service is an application service that runs the Informatica Analyst application in the Informatica domain. The Analyst Service manages the connections between service components and the users that log in to the Analyst tool. The Analyst Service connects to a Data Integration Service, a Model Repository Service, a Metadata Manager Service, and a Search Service.

The Analyst Service also specifies the following directories and database connections:

- Flat file cache directory
- Flat file network directory
- Business glossary export file directory
- Exception management audit database

When you configure the Analyst Service, connect it to a Data Integration Service to run profiles, scorecards, and mapping specifications. You can also connect the Analyst Service to a Data Integration Service that runs workflows that create Human tasks. Connect the Analyst Service to a Model Repository Service to identify a Model repository.

Connect the Analyst Service to a Metadata Manager Service to perform data lineage operations on scorecards in the Analyst tool. Connect the Analyst Service to a Search Service to manage search operations in the Analyst tool.

Specify a flat file cache directory to store temporary data from flat files that you upload. Specify a business glossary directory to stores temporary files that you export from the Business Glossary.

Specify a flat file network directory so that you can import flat files using the network path option in the Analyst tool.

Specify a database as the exception management audit database to write an audit trail for all Human task operations to a single database. When you specify the database, specify the schema for the audit tables also. The database stores audit data for all work that users perform on Human task instances in the Analyst tool

that the Analyst Service manages. If you do not specify a database and schema, the Analyst Service writes audit data for each Human task instance to the database that stores the Human task data.

Content Management Service

The Content Management Service is an application service that manages reference data. The service provides reference data properties to the Data Integration Service and to the Developer tool. The service also generates mapplets from rule specifications. You can create rule specifications and generate mapplets from rule specifications in the Analyst tool.

The Content Management Service must be available when you use the following resources:

Address reference data

The Content Management Service manages configuration information for address reference data. The Data Integration Service applies the configuration information when it runs a mapping that reads the address reference data.

Identity population files

The Content Management Service manages the list of the population files on the node. The Data Integration Service applies the population configuration when it runs a mapping that reads the population files.

Probabilistic model files and classifier model files

The Content Management Service stores the locations of any probabilistic model file and classifier model file on the node. The Content Management Service also manages the compilation status of each model.

Reference tables

The Content Management Service identifies the database that stores data values for the reference table objects in the associated Model repository.

Rule specifications

The Content Management Service generates mapplets from rule specifications. The Analyst Service selects a Content Management Service to generate the mapplets. The Analyst tool uses the Model Repository Service configuration to select the Content Management Service.

Data Integration Service

The Data Integration Service is an application service that performs data integration tasks for Informatica Analyst, Informatica Developer, and external clients. Data integration tasks include previewing data and running profiles, SQL data services, web services, and mappings.

When you start a command from the command line or an external client to run SQL data services and mappings in an application, the command sends the request to the Data Integration Service.

You can configure the Data Integration Service to run on the following domain objects:

On nodes

If your license includes high availability, you can configure the service to run on multiple nodes. By default, the service runs on the primary node. If the primary node is not available, the service runs on a back-up node. If the service process fails or the node becomes unavailable, the service fails over to another node. If your license does not include high availability, you can configure the service to run on one node.

On a grid

If your license includes grid, you can configure the service to run on a grid. A grid is an alias assigned to a group of nodes. The Data Integration Service dispatches jobs to available nodes assigned to the grid. When the Data Integration Service runs on a grid, the service remains available if a service process fails or a node becomes unavailable.

Email Service

The Email Service is a system service that manages email notifications for business glossaries, scorecards, and workflows.

Enable the Email Service to allow users to configure email notifications.

The Email Service sends emails for the following notifications:

- Business glossary notifications.
- Scorecard notifications.
- Workflow notifications. Workflow notifications include emails sent from Human tasks and Notification tasks in workflows that the Data Integration Service runs.

You can configure the service to run on multiple nodes. Designate the primary node to run the service. All other nodes are backup nodes for the service. If the primary node is not available, the service runs on a backup node.

Metadata Access Service

The Metadata Access Service is a user-managed service that allows the Developer tool to access Hadoop connection information to import and preview metadata.

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. Google Cloud Storage connection uses Metadata Access Service to import metadata from files in Google Cloud Storage. Create and configure a Metadata Access Service before you create Google Cloud Storage, HBase, HDFS, Hive, and MapR-DB connections.

Metadata Manager Service

The Metadata Manager Service is an application service that runs the Metadata Manager application and manages connections between the Metadata Manager components.

Use Metadata Manager to browse and analyze metadata from disparate source repositories. You can load, browse, and analyze metadata from application, business intelligence, data integration, data modelling, and relational metadata sources.

You can configure the Metadata Manager Service to run on only one node. The Metadata Manager Service is not a highly available service. However, you can run multiple Metadata Manager Services on the same node.

Model Repository Service

The Model Repository Service manages the Model repository. The Model repository stores metadata created by Informatica products in a relational database to enable collaboration among the products. Informatica

Developer, Informatica Analyst, Data Integration Service, and the Administrator tool store metadata in the Model repository.

You can configure a Model repository as a monitoring Model repository. You can then set up a monitoring Model Repository Service for the monitoring Model repository. The monitoring Model Repository Service monitors statistics for Data Integration Service jobs. You configure the monitoring Model Repository Service in the domain properties.

Create one Model Repository Service for each Model repository. When you create a Model Repository Service, you can create a Model repository or use an existing Model repository. You can run multiple Model Repository Services on the same node.

PowerCenter Integration Service

The PowerCenter Integration Service runs PowerCenter sessions and workflows. When you configure the PowerCenter Integration Service, you can specify where you want it to run:

- On a grid. When you configure the service to run on a grid, it can run on multiple nodes at a time. The PowerCenter Integration Service dispatches tasks to available nodes assigned to the grid. If you do not have the high availability option, the task fails if any service process or node becomes unavailable. If you have the high availability option, failover and recovery is available if a service process or node becomes unavailable.
- On nodes. If you have the high availability option, you can configure the service to run on multiple nodes. By default, it runs on the primary node. If the primary node is not available, it runs on a backup node. If the service process fails or the node becomes unavailable, the service fails over to another node. If you do not have the high availability option, you can configure the service to run on one node.

PowerCenter Repository Service

The PowerCenter Repository Service manages the PowerCenter repository. It retrieves, inserts, and updates metadata in the repository database tables. If the service process fails or the node becomes unavailable, then the service becomes unavailable.

If you have the high availability option, you can configure the service to run on primary and backup nodes. By default, the service process runs on the primary node. If the service process fails, a new process starts on the same node. If the node becomes unavailable, a service process starts on one of the backup nodes.

PowerExchange Listener Service

The PowerExchange Listener Service is an application service that manages the PowerExchange Listener. The PowerExchange Listener manages communication between a PowerCenter or PowerExchange client and a data source for bulk data movement and change data capture. The PowerCenter Integration Service connects to the PowerExchange Listener through the Listener Service. Use the Administrator tool to manage the service and view service logs.

If you have the PowerCenter high availability option, you can run the Listener Service on multiple nodes. If the Listener Service process fails on the primary node, it fails over to a backup node.

PowerExchange Logger Service

The Logger Service is an application service that manages the PowerExchange Logger for Linux, UNIX, and Windows. The PowerExchange Logger captures change data from a data source and writes the data to PowerExchange Logger log files. Use the Administrator tool to manage the service and view service logs.

If you have the PowerCenter high availability option, you can run the Logger Service on multiple nodes. If the Logger Service process fails on the primary node, it fails over to a backup node.

Resource Manager Service

The Resource Manager Service is a system service that manages computing resources in the domain and dispatches jobs to achieve optimal performance and scalability. The Resource Manager Service collects information about nodes with the compute role. The service matches job requirements with resource availability to identify the best compute node to run the job.

The Resource Manager Service communicates with compute nodes in a Data Integration Service grid. Enable the Resource Manager Service when you configure a Data Integration Service grid to run jobs in separate remote processes.

You can configure the service to run on multiple nodes. Designate the primary node to run the service. All other nodes are back-up nodes for the service. If the primary node is not available, the service runs on a back-up node.

SAP BW Service

The SAP BW Service listens for RFC requests from SAP NetWeaver BI and initiates workflows to extract from or load to SAP NetWeaver BI. The SAP BW Service is not highly available. You can configure it to run on one node.

Scheduler Service

The Scheduler Service is a system service that manages schedules for profiles, scorecards, deployed mappings, and deployed workflows.

Enable the Scheduler Service to create, manage, and run schedules.

You can configure the service to run on multiple nodes. Designate the primary node to run the service. All other nodes are back-up nodes for the service. If the primary node is not available, the service runs on a back-up node.

Web Services Hub

The Web Services Hub receives requests from web service clients and exposes PowerCenter workflows as services. The Web Services Hub does not run an associated service process. It runs within the Service Manager.

High Availability

High availability is an option that eliminates a single point of failure in a domain and provides minimal service interruption in the event of failure. High availability consists of the following components:

- Resilience. The ability of application services to tolerate transient network failures until either the resilience timeout expires or the external system failure is fixed.
- Failover. The migration of an application service or task to another node when the node running the service process becomes unavailable.

- **Recovery.** The automatic completion of tasks after a service is interrupted. Automatic recovery is available for PowerCenter Integration Service and PowerCenter Repository Service tasks. You can also manually recover PowerCenter Integration Service workflows and sessions. Manual recovery is not part of high availability.

Informatica Data Usage Policy

Informatica DiscoveryIQ is a product usage tool in the Informatica domain that sends routine reports on data usage and system statistics to Informatica.

Informatica DiscoveryIQ uploads data to Informatica 15 minutes after you install and configure the Informatica domain. Thereafter, the domain sends data to Informatica every few days. Data collection and upload is enabled by default. You can choose to not send any usage statistics to Informatica.

If the network where you install Informatica services needs a proxy server to communicate with the external network, configure the proxy details.

Informatica DiscoveryIQ enables Informatica to provide an environment health check after the analysis of system statistics and domain reports. You can receive best practices and recommendations from Informatica based on the reports to resolve issues on the domain. The usage statistics provide Informatica a proactive insight into product implementation.

Informatica DiscoveryIQ reports the following data to Informatica:

- Operating system details
- CPU information
- Informatica license key serial number
- Gateway information
- Domain options
- Node options
- Application service information

Configuring Informatica DiscoveryIQ Proxy Details

Configure proxy server details if the network on which you install Informatica services use a proxy server to communicate with the external network.

1. In the Administrator tool header area, click **Manage > DiscoveryIQ Proxy Details**.
2. Enter the domain, host name, and port number of the proxy server.
3. Enter the user name and password to connect to the proxy server.
4. Click **OK** to save the proxy server details.

Disabling Informatica Data Usage

You can disable the upload of usage data from the Informatica domain in the Administrator tool.

1. In the Administrator tool, click **Help > About**.
2. Click **Data Usage Policy**.
3. Clear **Enable Usage Collection**.

4. Click **OK**.

CHAPTER 2

Managing Your Account

This chapter includes the following topics:

- [Managing Your Account Overview, 28](#)
- [Password Management, 28](#)
- [Preferences, 29](#)
- [Informatica Network Credentials, 29](#)

Managing Your Account Overview

Manage your account to change your password or edit user preferences.

If you have a native user account, you can change your password at any time with the Change Password application. If someone else created your user account, change your password the first time you log in to the Administrator tool.

User preferences control the options that appear in the Administrator tool when you log in. User preferences do not affect the options that appear when another user logs in to the Administrator tool.

You can configure Informatica Network credentials for your account so that you can access the Informatica Knowledge Base from the Administrator tool.

Password Management

You can change the password through the Change Password application.

You can open the Change Password application from the Administrator tool or with the following URL:

`http://<fully qualified host name>:<port>/passwordchange/`

The Service Manager uses the user password associated with a worker node to authenticate the domain user. If you change a user password that is associated with one or more worker nodes, the Service Manager updates the password for each worker node. The Service Manager cannot update nodes that are not running. For nodes that are not running, the Service Manager updates the password when the nodes restart.

Note: For an LDAP user account, change the password in the LDAP directory service.

For a native user account, if you enable password complexity, use the following guidelines when you create or change a password:

- The length of the password must be at least eight characters.
- It must be a combination of an alphabet character, a numeric character and a non-alphanumeric character, such as:

! \ " # \$ % & ' () * + , - . / : ; < = > ? @ [] ^ _ ` { | } ~

When you use special characters in a password, the shell sometimes interprets them differently. For example, \$ is interpreted as a variable. In this case, use an escape character to escape the special character.

Changing Your Password

Change the password for a native user account at any time. For a user account created by someone else, change the password the first time you log in to the Administrator tool.

1. In the Administrator tool header area, click **Manage > Change Password**.
The Change Password application opens in a new browser window.
2. Enter the current password in the **Password** box, and the new password in the **New Password** and **Confirm Password** boxes.
3. Click **Update**.

Preferences

Your preferences determine the options that appear in the Administrator tool when you log in. Your preferences do not affect the options that appear when another user logs in to the Administrator tool.

The following table describes the options that you can configure for your preferences:

Option	Description
Subscribe for Alerts	Subscribes you to domain and service alerts. You must have a valid email address configured for your user account. Default is No.
Show Custom Properties	Displays custom properties in the contents panel when you click an object in the Navigator. You use custom properties to configure Informatica behavior for special cases or to increase performance. Hide the custom properties to avoid inadvertently changing the values. Use custom properties only if Informatica Global Customer Support instructs you to.

To edit your preferences, click **Manage > Preferences** in the Administrator tool header area.

Informatica Network Credentials

You can enter your Informatica Network credentials in the Administrator tool to access the Informatica Knowledge Base from the Administrator tool.

You can also view the search results for an error message in the Informatica Knowledge Base by clicking the error message code in the Administrator tool.

Enter Informatica Network Credentials

Enter your Informatica Network credentials to access the Informatica Knowledge Base from the Administrator tool.

1. Click **Manage > Support Portal Credentials**.
The **Informatica Network Login Credentials** window appears.
2. Enter your Informatica Network credentials and the customer project ID.
3. Click **OK**.

Searching Informatica Knowledge Base

You can search for terms in the Informatica Knowledge Base directly from the Administrator tool.

1. Click **Help > Search Knowledge Base**.
The **Search Knowledge Base** window appears.
2. Enter the term that you want to search in the text box.
3. Click **OK**.
The search results appear in a different browser window.

CHAPTER 3

Using Informatica Administrator

This chapter includes the following topics:

- [Using Informatica Administrator Overview, 31](#)
- [Log In to Informatica Administrator, 32](#)
- [Manage Tab, 33](#)
- [Manage Tab - Domain View, 34](#)
- [Manage Tab - Services and Nodes View, 38](#)
- [Manage Tab - Connections View, 46](#)
- [Manage Tab - Schedules View, 47](#)
- [Monitor Tab, 48](#)
- [Monitor Tab - Summary Statistics View, 49](#)
- [Monitor Tab - Execution Statistics View, 50](#)
- [Logs Tab, 56](#)
- [Reports Tab, 57](#)
- [Security Tab, 57](#)
- [Service States, 61](#)
- [Process States, 61](#)
- [Job States, 63](#)
- [Informatica Administrator Accessibility Overview, 63](#)

Using Informatica Administrator Overview

Informatica Administrator is the tool that you use to manage the Informatica domain and Informatica security.

Use the Administrator tool to complete the following types of tasks:

- Domain administrative tasks. Manage logs, domain objects, user permissions, and domain reports. Generate and upload node diagnostics. Monitor Data Integration Service jobs and applications. Domain objects include application services, nodes, grids, folders, database connections, operating system profiles, and licenses.
- Security administrative tasks. Manage users, groups, roles, and privileges.

The Administrator tool has the following tabs:

- **Manage.** View and edit the properties of the domain and objects within the domain.
- **Monitor.** View the status of profile jobs, scorecard jobs, preview jobs, mapping jobs, SQL data services, web services, and workflows for each Data Integration Service.
- **Logs.** View log events for the domain and services within the domain.
- **Reports.** Run a Web Services Report or License Management Report.
- **Security.** Manage users, groups, roles, and privileges.
- **Cloud.** View information about your Informatica Cloud® organization.

The Administrator tool has the following header items:

- **Log out.** Log out of the Administrator tool.
- **Manage.** Manage your account.
- **Help.** Access help for the current tab and determine the Informatica version.

Log In to Informatica Administrator

You must have a user account to log in to the Informatica Administrator web application.

If the Informatica domain runs on a network with Kerberos authentication, you must configure the browser to allow access to the Informatica web applications. In Microsoft Internet Explorer, Microsoft Edge, and Google Chrome, add the URL of the Informatica web application to the list of trusted sites. In Safari, add the certificate of the Informatica web application to the keychain. If you are using Chrome version 86.0.42x or later on Windows, you must also set the `AuthServerWhitelist` and `AuthNegotiateDelegateWhitelist` policies.

1. Start a Microsoft Internet Explorer or Google Chrome browser.
2. In the **Address** field, enter the URL for the Administrator tool:
 - If the Administrator tool is not configured to use a secure connection, enter the following URL:

```
http://<fully qualified hostname>:<http port>/administrator/
```

Note: You cannot access the Administrator tool with an unsecured connection if you have enabled the HTTP Strict Transport Security (HSTS) response header for the domain. For more information, see the *Informatica Security Guide*. You can use the HSTS option after you apply the Informatica 10.4.1.2 service pack.

- If the Administrator tool is configured to use a secure connection, enter the following URL:

```
https://<fully qualified hostname>:<https port>/administrator/
```

Host name and port in the URL represent the host name and port number of the master gateway node. If you configured secure communication for the domain, you must use HTTPS in the URL to ensure that you can access the Administrator tool.

If you use Kerberos authentication, the network uses single sign on. You do not need to log in to the Administrator tool with a user name and password.

3. If you do not use Kerberos authentication, enter the user name, password, and security domain for your user account, and then click **Login**.

The **Security Domain** field appears when the Informatica domain contains an LDAP security domain. If you do not know the security domain that your user account belongs to, contact the Informatica domain administrator.

Note: If this is the first time you log in with the user name and password provided by the domain administrator, change your password to maintain security.

Informatica Administrator URL

In the Administrator tool URL, `<host>:<port>` represents the host name of the master gateway node and the Administrator tool port number.

You configure the Administrator tool port when you define the domain. You can define the domain during installation or by running the `infasetup DefineDomain` command line program. If you enter the domain port instead of the Administrator tool port in the URL, the browser is directed to the Administrator tool port.

Note: If the domain fails over to a different master gateway node, the host name in the Administrator tool URL is equal to the host name of the elected master gateway node.

Troubleshooting the Login to Informatica Administrator

If the Informatica domain uses Kerberos authentication, you might encounter the following issues when logging in to the Administrator tool:

I cannot log in to the Administrator tool from the same machine where I created the domain gateway node.

After installation, if you cannot log in to the Administrator tool from the same machine where you created the domain gateway node, clear the browser cache. When you initially log in to the Administrator tool after installation, you can only log in with the Administrator user account created during installation. If a different user credential is stored in the browser cache, the login can fail.

A blank page appears after I log in to the Administrator tool.

If a blank page appears after you log in to the Administrator tool, verify that you enabled delegation for all user accounts with service principals used in the Informatica domain. To enable delegation, in the Microsoft Active Directory Service, set the **Trust this user for delegation to any service (Kerberos only)** option for each user account that you set an SPN.

Manage Tab

On the **Manage** tab, you can view and manage the domain and the objects that it contains.

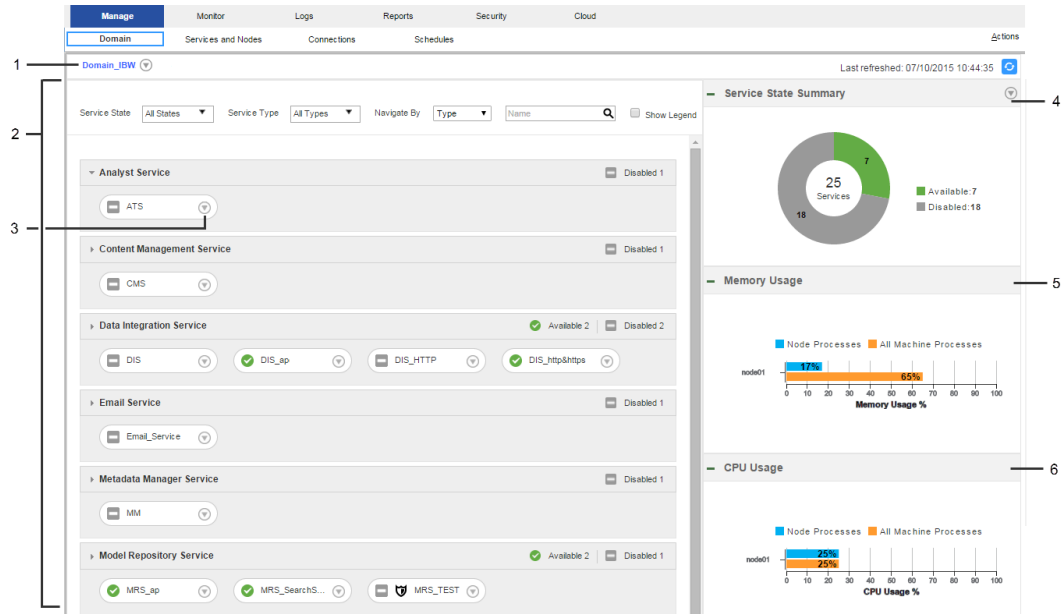
The contents that appear and the tasks you can complete on the **Manage** tab vary based on the view that you select. You can select the following views:

- **Domain.** View and manage domain status, resource consumption, and events.
- **Services and Nodes.** View and manage application services and nodes.
- **Connections.** View and manage connections.
- **Schedules.** View and manage schedules for profiles, scorecards, deployed mappings, and deployed workflows.

Manage Tab - Domain View

The **Domain** view displays an overview of the domain and its contents. You can use the domain view to monitor the domain status, resource consumption, and events. You can also perform domain actions, such as enabling and disabling services.

The following image shows the **Domain** view on the **Manage** tab:



1. Domain Actions menu
2. Contents panel
3. Object Actions menu
4. Service State Summary
5. Memory usage indicator
6. CPU usage indicator

The **Domain** view has the following components:

Domain Actions menu

Use the Domain Actions menu to view more information about the domain or shut down the domain.

Use the Domain Actions menu to perform the following tasks:

- View Properties. Open the Services and Nodes view and display the properties for the domain.
- View History. Open the History view and display domain events from the last day.
- View Logs. Open the Logs tab and display Service Manager log events from the last day.
- View Command History. Open the Command History panel and display the 50 most recent service lifecycle commands that were issued from the Administrator tool.
- Shut Down Domain.

Contents panel

Displays domain objects and their states. Domain objects include services, nodes, and grids.

The following table describes the methods that you can use to view objects in the contents panel:

Method	Description
Service State	Filter services by the following states: <ul style="list-style-type: none"> - All states - Available - Unavailable - Disabled
Service Type	Filter some or all services in the domain.
Navigate by	Group objects by node, type, or folder.
Search	Search for an object by name. You can use an asterisk (*) as a wildcard character in this field.
Show Legend	View a list of state icons and descriptions.

Object Actions menus

Objects in the contents panel have Actions menus. Use the Actions menus to view information about domain objects or perform common tasks. The information you can view and the tasks that you can perform vary depending on which object you select.

Use the service Actions menu to perform the following tasks:

- View Properties. Open the Services and Nodes view and display the properties for the service.
- View History. Open the History view and display service events from the last day.
- View Logs. Open the Logs tab and display service log events from the last day.
- View Dependencies. Open the Dependency graph and display direct dependencies for the service.
- Recycle Service.
- Enable or Disable Service.

Use the node Actions menu to perform the following tasks:

- View Properties. Open the Services and Nodes view and display properties for the node.
- View History. Open the History view and display node events from the last day.
- View Dependencies. Open the Dependency graph and display direct dependencies for the node.
- Shut Down Node

Use the grid Actions menu to perform the following tasks:

- View Properties. Open the Services and Nodes view and display properties for the grid.
- View Dependencies. Open the Dependency graph and display direct dependencies for the grid.

Service State Summary

Doughnut chart that shows the number and state of the services in the domain. When you click a state in the chart, the contents panel displays services with that state.

Resource usage indicators

Bar charts that compare resource usage of a domain process to resource usage of all processes on the machine. The Domain view contains a memory usage indicator and a CPU usage indicator.

Manage tab Actions menu

Access help for the Domain view.

Details Panel

When you select a domain object, the **Details** panel displays information about the object. The information that you can view varies depending on which object you select.

The following table describes the details that display depending on the object that you select in the contents panel:


Object	Details Panel Content
Node	Node name and state. Click the node name to view the node properties.
Service	The Details panel displays the following content for a service: <ul style="list-style-type: none">- Service name and state. Click the service name to view the service properties.- Node on which the service process runs. Click the node name to view the node properties.- State of the node on which the service process runs.- State of the service process.
Service running on a grid	The Details panel displays the following content for a service that runs on a grid: <ul style="list-style-type: none">- Service name and state. Click the service name to view the service properties.- Nodes in the grid. Click a node name to view the node properties.- State of the nodes on which the service processes run.- State of the service processes.
Service in high availability mode	The Details panel displays the following content for a service that is highly available: <ul style="list-style-type: none">- Service name and state. Click the service name to view the service properties.- Nodes on which the service is configured to run. Click a node name to view the node properties.- State of the nodes on which the service processes run.- State of the service process on the nodes.
Grid	The Details panel displays the following content for a grid: <ul style="list-style-type: none">- Grid name and state. Click the grid name to view the grid properties.- Nodes in the grid. Click a node name to view the node properties.- State of nodes running in the grid.

Resource Usage Indicators

The resource usage indicators are bar charts and graphs that compare resource usage for a domain process to resource usage of all processes on the machine. Select a domain process to compare with all processes. You can view current usage statistics or statistics for the previous 60 minutes.

You can view usage statistics for memory and CPU. Choose whether to view current statistics or to view graphs of usage for the past 60 minutes. Click the selection arrow and choose **Current** or **Last Hour Trend**.

The following image shows the current resource usage in a domain that contains one node:



The information that the graphs display varies based on which domain object you select.

The following table describes the information that you can view when you select the domain or a domain object:

Domain Object	Usage Indicator Content
Domain	<p>The usage indicators display the following content:</p> <ul style="list-style-type: none"> - Nodes in the domain. - Resource usage of all processes running on each node in the domain. - Resource usage of all processes running on the machine.
Node	<p>The usage indicators display the following content:</p> <ul style="list-style-type: none"> - The node. - Resource usage of processes running on the node. - Resource usage of all processes running on the machine.
Service	<p>The usage indicators display the following content:</p> <ul style="list-style-type: none"> - The node on which the service process runs. - Resource usage of the service process that is running on the node. - Resource usage of all processes running on the machine.
Service in high availability mode	<p>The usage indicators display the following content:</p> <ul style="list-style-type: none"> - The node on which the service process is running. - Resource usage of the service process that is running on the node. - Resource usage of all processes running on the machine.

Domain Object	Usage Indicator Content
Service running on a grid	The usage indicators display the following content: <ul style="list-style-type: none"> - All nodes on which the service process runs. - Resource usage of the service process that is running on each node. - Resource usage of all processes running on the machine.
Grid	The usage indicators display the following content: <ul style="list-style-type: none"> - All available nodes in the grid. - Resource usage of all processes running on each node in the domain. - Resource usage of all processes running on the machine.

If a **See More...** link appears in the indicator, you can click it to view the complete list of nodes in the domain. You can sort the list by node name, process usage on the nodes, or process usage on the machine. You can also search the list for a particular node.

Manage Tab - Services and Nodes View

The **Services and Nodes** view shows all application services and nodes defined in the domain.

The **Services and Nodes** view has the following components:

Navigator

Appears in the left pane of the **Manage** tab. The Navigator displays the following types of objects:

- Domain. You can view one domain, which is the highest object in the Navigator hierarchy.
- Folders. Use folders to organize domain objects in the Navigator. Select a folder to view information about the folder and the objects in the folder.
- Application services. An application service represents server-based functionality. Select an application service to view information about the service and its processes.
- System services. A system service is an application service that can have a single instance in the domain. Select a system service to view information about the service and its processes.
- Nodes. A node represents a machine in the domain. You configure service processes to run on nodes with the service role.
- Grids. Create a grid to run the Data Integration Service or PowerCenter Integration Service on multiple nodes. Select a grid to view nodes assigned to the grid.
- Licenses. Create a license on the **Manage** tab based on a license key file provided by Informatica. Select a license to view services assigned to the license.

You can search for nodes, services, and grids in the Navigator.

Contents panel

Appears in the right pane of the **Manage** tab and displays information about the domain or domain object that you select in the Navigator.

Actions menu in the Navigator

When you select the domain in the Navigator, you can create a folder, service, node, grid, or license.

When you select a domain object in the Navigator, you can delete the object, move it to a folder, or refresh the object.

Actions menu on the Manage tab

When you select the domain in the Navigator, you can shut down the domain or view logs for the domain.

When you select a node in the Navigator, you can remove a node association, recalculate the CPU profile benchmark, or shut down the node.

When you select a service in the Navigator, you can recycle or disable the service and configure properties for the service.

When you select a license in the Navigator, you can add an incremental key to the license.

Navigator Search

You can search for and filter nodes, application services, and grids in the Navigator.

You can perform the following tasks in the Navigator search section:

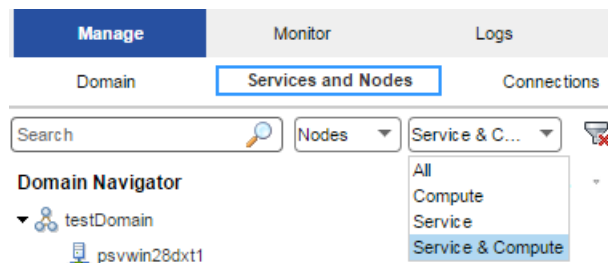
Search by object name.

In the search text box, enter the name or partial name of the object to search for. The Navigator displays the search results.

Filter by object type.

Click **Filters**, and then select the object type to filter by. If you filter by nodes, you can filter further by node role. If you filter by services, you can filter further by service type. The Navigator displays the filtered results.

The following image displays the Navigator search section filtered by nodes with the service and compute roles:



Reset filters.

Click **Reset Filters** to clear any filters or entered search text.

Domain

You can view one domain in the **Services and Nodes** view on the **Manage** tab. It is the highest object in the Navigator hierarchy.

When you select the domain in the Navigator, the contents panel shows the following views and buttons:

- **Properties** view. View or edit domain resilience properties.
- **Resources** view. View available resources for each node in the domain.
- **Permissions** view. View or edit group and user permissions on the domain.

- **Diagnostics** view. View node diagnostics, generate and upload node diagnostics to Customer Support Manager, or edit customer portal login details.
- **Plug-ins** view. View plug-ins registered in the domain.
- **View Domain Logs** button. View logs for the domain and services in the domain.

In the **Actions** menu in the Navigator, you can add a folder, node, grid, application service, or license to the domain.

In the **Actions** menu on the **Manage** tab, you can shut down the domain, view logs, or access help for the current view.

Folders

You can use folders in the domain to organize objects and to manage security.

Folders can contain nodes, services, grids, licenses, and other folders.

When you select a folder in the Navigator, the Navigator opens to display the objects in the folder. The contents panel displays the following information:

- **Properties** view. Displays the name and description of the folder.
- **Permissions** view. View or edit group and user permissions on the folder.

In the **Actions** menu in the Navigator, you can delete the folder, move the folder into another folder, refresh the contents on the **Manage** tab, or access help for the current tab.

Note: The System_Services folder is created for you when you create the domain, and contains all of the system services. A system service is an application service that can have a single instance in the domain. You cannot delete, move, or edit the properties or contents of the System_Services folder.

Application Services

Application services are a group of services that represent Informatica server-based functionality.

In the **Services and Nodes** view on the **Manage** tab, you can create and manage the following application services:

Analyst Service

Runs Informatica Analyst in the Informatica domain. The Analyst Service manages the connections between service components and the users that log in to Informatica Analyst.

The Analyst Service connects to a Data Integration Service, a Model Repository Service, a Metadata Manager Service, and a Search Service. The Analyst Service also specifies a flat file cache directory and flat file network directory, a directory for business glossary export files, and an exception management audit database.

You can create and recycle the Analyst Service in the Informatica domain to access the Analyst tool. You can launch the Analyst tool from the Administrator tool.

When you select an Analyst Service in the Navigator, the contents panel displays the following information:

- **Properties** view. View the state of the service and the URL for the Analyst Service. Manage general, model repository, data integration, human task, metadata manager, flat file cache, flat file network, business glossary export, logging, and custom properties.
- **Processes** view. View the state of the service process on each node. View and edit service process properties on each assigned node.

- **Permissions** view. View or edit the group and user permissions on the Analyst Service.
- **Actions** menu. Manage the service and repository contents.

Content Management Service

Manages reference data and compiles rule specifications into mapplets.

When you select a Content Management Service in the Navigator, the contents panel displays the following information:

- **Properties** view. View the state of the service. Manage general, master, data integration, model repository, reference table data, temporary file, logging, and custom properties.
- **Processes** view. View the state of the service process on each node. View and edit service process properties on each assigned node.
- **Permissions** view. View or edit the group and user permissions on the Content Management Service.
- **Actions** menu. Manage the service.

Data Integration Service

Completes data integration tasks for Informatica Analyst, Informatica Developer, and external clients. When you preview or run data profiles, SQL data services, and mappings in Informatica Analyst or Informatica Developer, the application sends requests to the Data Integration Service to perform the data integration tasks. When you start a command from the command line or an external client to run SQL data services and mappings in an application, the command sends the request to the Data Integration Service.

When you select a Data Integration Service in the Navigator, the contents panel displays the following information:

- **Properties** view. View the state of the service. Manage general, model repository, logging, logical data object and virtual table cache, profiling, data object cache, and custom properties. Set the default deployment option.
- **Processes** view. View the state of the service process on each node. View and edit service process properties on each assigned node.
- **Applications** view. Start and stop applications and SQL data services. Back up applications. Manage application properties.
- **Actions** menu. Manage the service and repository contents.

Metadata Access Service

Allows the Developer tool to access Hadoop connection information to import and preview metadata. When importing a data object, the Developer tool sends a request to the Metadata Access Service.

When you select a Metadata Access Service in the Navigator, the contents panel displays the following information:

- **Properties** view. View the state of the service. Manage general, execution, HTTP configuration and custom properties.
- **Processes** view. View the state of the service process on each node. View and edit service process properties on each assigned node.
- **Permissions** view. View or edit the group and user permissions on the Metadata Access Service.
- **Actions** menu. Manage the service.

Metadata Manager Service

Runs the Metadata Manager application and manages connections between the Metadata Manager components.

When you select a Metadata Manager Service in the Navigator, the contents panel displays the following information:

- **Properties** view. View the state of the service and the URL of the Metadata Manager Service instance. View or edit Metadata Manager properties.
- **Associated Services** view. View and configure the Integration Service associated with the Metadata Manager Service.
- **Permissions** view. View or edit the group and user permissions on the Metadata Manager Service.
- **Actions** menu. Manage the service and repository contents.

Model Repository Service

Manages the Model repository. The Model repository stores metadata created by Informatica products, such as Informatica Developer, Informatica Analyst, the Data Integration Service, and Informatica Administrator. The Model repository enables collaboration among the products.

When you select a Model Repository Service in the Navigator, the contents panel displays the following information:

- **Properties** view. View the state of the service. Manage general, repository database, search, and custom properties.
- **Processes** view. View the state of the service process on each node. View and edit service process properties on each assigned node.
- **Actions** menu. Manage the service and repository contents.

You can configure a Model repository as a monitoring Model repository. You can configure a monitoring Model Repository Service at the domain level to monitor multiple Data Integration Service and the objects that the Data Integration Service runs.

PowerCenter Integration Service

Runs PowerCenter sessions and workflows. Select a PowerCenter Integration Service in the Navigator to access information about the service.

When you select a PowerCenter Integration Service in the Navigator, the contents panel displays the following information:

- **Properties** view. View the state of the service. View or edit Integration Service properties.
- **Associated Repository** view. View or edit the repository associated with the Integration Service.
- **Processes** view. View the state of the service process on each node. View or edit the service process properties on each assigned node.
- **Permissions** view. View or edit group and user permissions on the Integration Service.
- **Actions** menu. Manage the service.

PowerCenter Repository Service

Manages the PowerCenter repository. It retrieves, inserts, and updates metadata in the repository database tables. Select a PowerCenter Repository Service in the Navigator to access information about the service.

When you select a PowerCenter Repository Service in the Navigator, the contents panel displays the following information:

- **Properties** view. View the state and operating mode of the service. Manage general and advanced properties, node assignments, and database properties.

- **Processes** view. View the state of the service process on each node. View and edit service process properties on each assigned node.
- **Connections and Locks** view. View and terminate repository connections and object locks.
- **Plug-ins** view. View and manage registered plug-ins.
- **Permissions** view. View or edit group and user permissions on the PowerCenter Repository Service.
- **Actions** menu. Manage the contents of the repository and perform other administrative tasks.

PowerExchange Listener Service

Runs the PowerExchange Listener.

When you select a Listener Service in the Navigator, the contents panel displays the following information:

- **Properties** view. View the state of the service and the URL of the PowerExchange Listener instance. View or edit Listener Service properties.
- **Actions** menu. Contains actions that you can perform on the Listener Service, such as viewing logs or enabling and disabling the service.

PowerExchange Logger Service

Runs the PowerExchange Logger for Linux, UNIX, and Windows.

When you select a Logger Service in the Navigator, the contents panel displays the following information:

- **Properties** view. View the state of the service and the URL of the PowerExchange Logger instance. View or edit Logger Service properties.
- **Actions** menu. Contains actions that you can perform on the Logger Service, such as viewing logs or enabling and disabling the service.

SAP BW Service

Listens for RFC requests from SAP BW and initiates workflows to extract from or load to SAP BW. Select an SAP BW Service in the Navigator to access properties and other information about the service.

When you select an SAP BW Service in the Navigator, the contents panel displays the following information:

- **Properties** view. View the state of the service. Manage general properties and node assignments.
- **Associated Integration Service** view. View or edit the Integration Service associated with the SAP BW Service.
- **Processes** view. View the state of the service process on each node. View or edit the directory of the BWParam parameter file.
- **Permissions** view. View or edit group and user permissions on the SAP BW Service.
- **Actions** menu. Manage the service.

Web Services Hub

A web service gateway for external clients. It processes SOAP requests from web service clients that want to access PowerCenter functionality through web services. Web service clients access the PowerCenter Integration Service and PowerCenter Repository Service through the Web Services Hub.

When you select a Web Services Hub in the Navigator, the contents panel displays the following information:

- **Properties** view. View the state of the service. View or edit Web Services Hub properties.

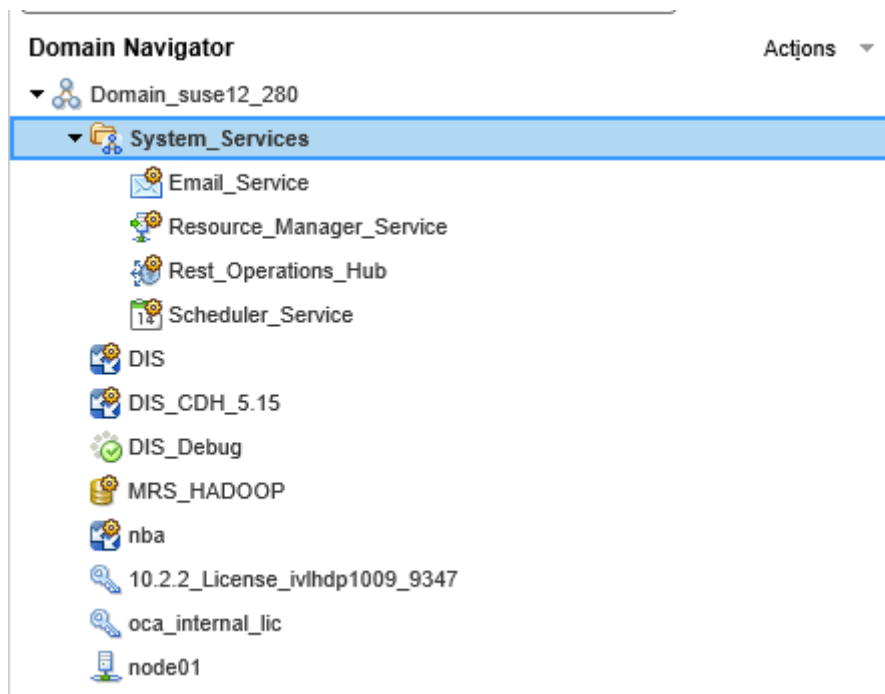
- **Associated Repository** view. View the PowerCenter Repository Services associated with the Web Services Hub.
- **Permissions** view. View or edit group and user permissions on the Web Services Hub.
- **Actions** menu. Manage the service.

System Services

A system service is an application service that can have a single instance in the domain. When you create the domain, the system services are created for you. You can enable, disable, and configure system services.

You can manage system services on the **Services and Nodes** view on the **Manage** tab.

The following image shows the System Services folder in the Domain Navigator:



You can manage the following system services:

Email Service

The Email Service sends email notifications for business glossaries, scorecards, and workflows. Enable the Email Service to allow users to configure email notifications.

When you select the Email Service in the System Services folder in the Navigator, the contents panel displays the following views and buttons:

- **Properties** view. View the state of the service. Manage the Model repository options and email server properties for the service.
- **Processes** view. View the state of the service process for each assigned node.
- **Permissions** view. View or edit the group and user permissions on the service.
- **Actions** menu. Manage the service.

Resource Manager Service

The Resource Manager Service manages computing resources in the domain and dispatches jobs to achieve optimal performance and scalability. The Resource Manager Service collects information about nodes with the compute role. The service matches job requirements with resource availability to identify the best compute node to run the job.

The Resource Manager Service communicates with compute nodes in a Data Integration Service grid. Enable the Resource Manager Service when you configure a Data Integration Service grid to run jobs in separate remote processes.

When you select the Resource Manager Service in the System Services folder in the Navigator, the contents panel displays the following views and buttons:

- **Properties** view. View the state of the service. Manage the log level and the primary and backup nodes for the service.
- **Processes** view. View the state of the service process for each assigned node.
- **Permissions** view. View or edit the group and user permissions on the service.
- **Actions** menu. Manage the service.

REST Operations Hub Service

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

Scheduler Service

The Scheduler Service manages schedules for profiles, scorecards, deployed mappings, and deployed workflows. Enable the Scheduler Service to create, manage, and run schedules.

When you select the Scheduler Service in the System Services folder in the Navigator, the contents panel displays the following views and buttons:

- **Properties** view. View the state of the service. Manage the log level, primary and backup nodes, and Model Repository Service options for the service.
- **Processes** view. View the state of the service process for each assigned node and configure process properties.
- **Permissions** view. View or edit the group and user permissions on the service.
- **Actions** menu. Manage the service.

Nodes

A node is a logical representation of a physical machine in the domain. On the Services and Nodes view on the Manage tab, you assign resources to nodes and configure service processes to run on nodes that have the service role.

When you select a node in the Navigator, the contents panel displays the following information:

- **Properties** view. View the status of the node. View or edit node properties, such as the repository backup directory or range of port numbers for the processes that run on the node.
- **Processes** view. View the status of processes configured to run on the node. Service processes run on nodes that have the service role.
- **Resources** view. View or edit resources assigned to the node.
- **Permissions** view. View or edit group and user permissions on the node.

In the **Actions** menu in the Navigator, you can delete the node, move the node to a folder, refresh the contents on the **Manage** tab, or access help on the current tab.

In the **Actions** menu on the **Manage** tab, you can remove the node association, recalculate the CPU profile benchmark, or shut down the node.

Grids

A grid is an alias assigned to a group of nodes that run PowerCenter Integration Service or Data Integration Service jobs.

When you run a job on a grid, the Integration Service distributes the processing across multiple nodes in the grid. For example, when you run a profile on a grid, the Data Integration Service splits the work into multiple jobs and assigns each job to a node in the grid. You assign nodes to the grid in the **Services and Nodes** view on the **Manage** tab.

When you select a grid in the Navigator, the contents panel displays the following information:

- **Properties** view. View or edit node assignments to a grid.
- **Permissions** view. View or edit group and user permissions on the grid.

In the **Actions** menu in the Navigator, you can delete the grid, move the grid to a folder, refresh the contents on the **Manage** tab, or access help for the current tab.

Licenses

You create a license object on the **Manage** tab based on a license key file provided by Informatica.

After you create the license, you can assign services to the license.

When you select a license in the Navigator, the contents panel displays the following information:

- **Properties** view. View license properties, such as supported platforms, repositories, and licensed options. You can also edit the license description.
- **Assigned Services** view. View or edit the services assigned to the license.
- **Options** view. View the licensed PowerCenter options.
- **Permissions** view. View or edit user permissions on the license.

In the **Actions** menu in the Navigator, you can delete the license, move the license to a folder, refresh the contents on the **Manage** tab, or access help on the current tab.

In the **Actions** menu on the **Manage** tab, you can add an incremental key to a license.

Manage Tab - Connections View

The **Connections** view shows the domain and all connections in the domain.

The **Connections** view has the following components:

Navigator

Displays the domain and the connections in the domain.

Actions menu in the Navigator

When you select the domain in the Navigator, you can create a connection.

When you select a connection in the Navigator, you can delete the connection.

Contents panel

Displays information about the domain or the connection that you select in the Navigator.

When you select the domain in the Navigator, the contents panel shows all connections in the domain. In the contents panel, you can filter or sort connections, or search for specific connections.

When you select a connection in the Navigator, the contents panel displays information about the connection. The tasks that you can complete for the connection vary depending on which of the following views you select:

- **Properties** view. View or edit connection properties.
- **Pooling** view. View or edit pooling properties for the connection.
- **Permissions** view. View or edit group or user permissions on the connection.

Actions menu on the Manage tab

When you select a connection in the Navigator, you can test the connection.

Manage Tab - Schedules View

Use the **Schedules** view to view and manage schedules for deployed mappings and workflows.

The **Schedules** view has the following components:

Navigator

Displays the domain and schedules in the domain.

All Schedules view

When you select the domain in the Navigator, the **All Schedules** view displays the name, status, and description of the schedules in the domain.

Properties view

When you select a schedule in the Navigator, the **Properties** view displays details about the schedule.

Scheduled Jobs view

When you select the domain in the Navigator, the **Scheduled Jobs** view displays details about the scheduled jobs in the domain.

When you select a schedule in the Navigator, the **Scheduled Jobs** view displays details about the jobs that are running on that schedule.

Actions menu on the Manage tab

The actions that you can perform in the Actions menu vary based on which view you select.

You can perform the following actions using the Actions menu:

- Create a schedule
- Edit a schedule
- Delete a schedule
- Pause a schedule
- Resume a schedule

- Refresh
- Unschedule an object
- Reset filters
- Access help

Monitor Tab

On the **Monitor** tab, you can monitor Data Integration Services and objects that the Data Integration Services run.

After you configure the monitoring Model repository on the domain, you can view monitoring statistics in the Administrator tool.

You can monitor the following objects on the **Monitor** tab:

- Ad hoc jobs. Includes profiles, enterprise data profiles, mappings, scorecards, reference tables, and previews.
- Applications. Includes deployed mappings, logical data objects, SQL data services, web services, and workflows.
- SQL data service connections.
- Requests. Includes SQL data service and web service requests.


You can select the following views:

- **Summary Statistics** view. View graphical information about object state, distribution, and resource usage across the Data Integration Services.
- **Execution Statistics** view. View properties, statistics, and reports for the objects that the Data Integration Service runs.

Monitor Tab - Summary Statistics View

The **Summary Statistics** view displays historical statistics about Data Integration Services and the jobs that they run.

The following image shows the **Summary Statistics** view:



Use the **Summary Statistics** view to view graphical information about object state, distribution, and resource usage.

The **Summary Statistics** view contains the following components:

Time line

Specify a source and time range for which you want to view statistics.

Ad Hoc Jobs panel

View doughnut or tabular charts for jobs for the selected source and time range. Select sections of the doughnuts to filter by job type, state, and Data Integration Service.

Deployed Mappings and Workflows panel

View doughnut or tabular charts for deployed mappings and workflows for the selected source and time range. Select sections of the doughnuts to filter by type, state, and Data Integration Service.

Requests and Connections panel

View doughnut or tabular charts for Data Integration Service jobs for the selected source and time range. Select sections of the doughnuts to filter by object type, state, and Data Integration Service.

Resource Usage panel

View CPU and memory usage of all of the Data Integration Service processes that are running in the domain or on a node in the domain.

Monitor Tab - Execution Statistics View

On the **Execution Statistics** view on the **Monitor** tab, you can monitor Data Integration Services and objects that the Data Integration Services run.

The following image shows the **Execution Statistics** view on the **Monitor** tab:

Name	Type	State	Job ID	Started By	Start Time	Elapsed Time	End Time
Read_flat	Preview	Canceled	wapFH242Ee...	Administrator	09/30/2015 02:58:10	00:00:02	09/30/2015 02:58:12
Map_flat	Mapping	Canceled	tf0F32d2EeW...	Administrator	09/30/2015 02:57:50	00:00:09	09/30/2015 02:58:00
Map_flat	Mapping	Completed	NkBUEmTE...	Administrator	09/30/2015 02:11:19	00:00:08	09/30/2015 02:11:28

General Properties	
Name	Map_flat
Type	Mapping
Started By	Administrator
User Security Domain	Native
Start Time	09/30/2015 02:57:50
Elapsed Time	00:00:09
End Time	09/30/2015 02:58:00

1. Navigator
2. Folders
3. Actions menu
4. Contents panel
5. Details panel
6. Views in the details panel

When you select an object in the Navigator, you can view details about the object and monitor the object.

You can select the following types of objects in the Navigator in the **Execution Statistics** view:

Domain

View the states and properties of Data Integration Services in the domain.

Data Integration Service

View general properties about the Data Integration Service, and view statistics about objects that the Data Integration Service runs.

Folder

View a list of objects in the folder. The folder is a logical grouping of objects. When you select a folder, a list of objects appears in the contents panel. The contents panel shows multiple columns that show properties about each object. You can configure the columns that appear in the contents panel.

The following table shows the folders that appear in the Navigator:

Folder	Location
Ad Hoc Jobs	Appears under the Data Integration Service.
Deployed Mapping Jobs	Appears under the corresponding application.
Logical Data Objects	Appears under the corresponding application.
SQL Data Services	Appears under the corresponding application.
Web Services	Appears under the corresponding application.
Workflows	Appears under the corresponding application.

Views in the Execution Statistics View

When you select an integration object in the Navigator or an object link in the contents panel of the **Execution Statistics** view, multiple views of information appear in the contents panel. The views show information about the selected object, such as properties, run-time statistics, and run-time reports.

Depending on the type of object you select in the Navigator, the contents panel might display the following views:

Properties view

Shows general properties and run-time statistics about the selected object. General properties might include the name and description of the object. Statistics vary based on the selected object type.

Statistics view

Shows historical statistics about jobs in an application, or jobs that the Data Integration Service ran. For example, when you select an application, you can view the number of deployed mapping jobs that failed in the last four hours.

Reports view

Shows reports for the selected object. The reports contain key metrics for the object. For example, you can view reports to determine the longest running jobs on a Data Integration Service during a particular time period.

Summary Statistics view

Shows throughput and resource usage statistics for ad hoc mapping jobs, deployed mapping jobs, or mappings in a workflow.

Detailed Statistics view

Shows graphs of the throughput and resource usage for ad hoc mapping jobs, deployed mapping jobs, or mappings in a workflow.

Historical Statistics view

Shows averaged data from multiple runs for a specific job. For example, you can view the minimum, maximum, and average duration of the mapping job. You can view the average amount of CPU that the job consumes when it runs.

Connections view

Shows connections defined for the selected object. You can view statistics about each connection, such as the number of closed, aborted, and total connections.

Requests view

Shows details about requests. There are two types of requests: SQL queries and web service requests. Users can use a third-party client tool to run SQL queries against the virtual tables in an SQL data service. Users can use a web service client to run web service requests against a web service. Each web service request runs a web service operation.

A request is a web services request or an SQL query that a user runs against a virtual table in an SQL data service.

Virtual Tables view

Shows virtual tables defined in an SQL data service. You can also view properties and cache refresh details for each virtual table.

Operations view

Shows the operations defined for the web service.

Statistics in the Execution Statistics View

When you select a Data Integration Service or an application in the Navigator in the **Execution Statistics** view, the **Statistics** section in the **Properties** view shows statistics for jobs that run on the Data Integration Service.

The following table describes the types of jobs and statistics that you can view:

Object Type	Statistics
Ad Hoc Jobs	Displays the following statistics about ad hoc jobs: <ul style="list-style-type: none">- Total. Total number of jobs.- Failed. Number of failed jobs.- Aborted. Number of aborted jobs. The Data Integration Service was recycled, or disabled in the abort mode when the job was running.- Completed. Number of completed jobs.- Canceled. Number of canceled jobs.
Applications	Displays the following statistics about application: <ul style="list-style-type: none">- Total. Total number of applications.- Running. Number of running applications.- Failed. Number of failed applications.- Stopped. Number of stopped applications.- Disabled. Number of disabled applications.
Deployed Mapping Jobs	Displays the following statistics about deployed mapping jobs: <ul style="list-style-type: none">- Total. Total number of deployed mappings.- Failed. Number of failed mapping jobs.- Aborted. Number of aborted mapping jobs.- Completed. Number of completed mapping jobs.- Canceled. Number of canceled mapping jobs.

Object Type	Statistics
Connections	Displays the following statistics about SQL data service connections to virtual databases: <ul style="list-style-type: none"> - Total. Total number of connections. - Closed. Number of database connections on which SQL data service requests have previously run. - Aborted. Number of connections that were manually aborted, or that were aborted when the Data Integration Service was recycled or disabled in abort mode.
Requests	Displays the following statistics about SQL data service and web service requests: <ul style="list-style-type: none"> - Total. Total number of requests. - Completed. Number of completed requests. - Aborted. Requests that were aborted when the Data Integration Service was recycled, or disabled in abort mode. - Failed. Number of failed requests.
Workflows	Displays the following statistics about workflows: <ul style="list-style-type: none"> - Total. Total number of workflow instances. - Completed. Number of completed workflow instances. - Canceled. Number of canceled workflow instances. - Aborted. Number of aborted workflow instances. - Failed. Number of failed workflow instances.

Reports on the Execution Statistics View

You can view monitoring reports in the **Reports** view of the **Execution Statistics** view. The **Reports** view appears when you select the appropriate object in the Navigator. You can view reports to monitor objects deployed to a Data Integration Service, such as ad hoc jobs, web services, SQL data services, and workflows.

The reports that appear in the **Reports view** vary depending on which object you select. To view reports in the **Reports** view, you must configure them in **Monitor tab Actions > Report and Statistic Settings**. By default, no reports appear in the **Reports** view.

You can view the following monitoring reports:

Longest Duration Ad Hoc Jobs

Shows ad hoc jobs that ran the longest during the specified time period. The report shows the job name, ID, type, state, and duration. You can view this report in the **Reports** view when you monitor a Data Integration Service in the **Monitor** tab.

Longest Duration Mapping Jobs

Shows mapping jobs that ran the longest during the specified time period. The report shows the job name, state, ID, and duration. You can view this report in the **Reports** view when you monitor a Data Integration Service in the **Monitor** tab.

Longest Duration Profile Jobs

Shows profile jobs that ran the longest during the specified time period. The report shows the job name, state, ID, and duration. You can view this report in the **Reports** view when you monitor a Data Integration Service in the **Monitor** tab.

Longest Duration Reference Table Jobs

Shows reference table process jobs that ran the longest during the specified time period. Reference table jobs are jobs where you export or import reference table data. The report shows the job name, state, ID, and duration. You can view this report in the **Reports** view when you monitor a Data Integration Service in the **Monitor** tab.

Longest Duration Scorecard Jobs

Shows scorecard jobs that ran the longest during the specified time period. The report shows the job name, state, ID, and duration. You can view this report in the **Reports** view when you monitor a Data Integration Service in the **Monitor** tab.

Longest Duration SQL Data Service Connections

Shows SQL data service connections that were open the longest during the specified time period. The report shows the connection ID, SQL data service, connection state, and duration. You can view this report in the **Reports** view when you monitor a Data Integration Service, an SQL data service, or an application in the **Monitor** tab.

Longest Duration SQL Data Service Requests

Shows SQL data service requests that ran the longest during the specified time period. The report shows the request ID, SQL data service, request state, and duration. You can view this report in the **Reports** view when you monitor a Data Integration Service, an SQL data service, or an application in the **Monitor** tab.

Longest Duration Web Service Requests

Shows web service requests that were open the longest during the specified time period. The report shows the request ID, web service operation, request state, and duration. You can view this report in the **Reports** view when you monitor a Data Integration Service, a web service, or an application in the **Monitor** tab.

Longest Duration Workflows

Shows all workflows that were running the longest during the specified time period. The report shows the workflow name, state, instance ID, and duration. You can view this report in the **Reports** view when you monitor a Data Integration Service or an application in the **Monitor** tab.

Longest Duration Workflows Excluding Human Tasks

Shows workflows that do not include a Human task that were running the longest during the specified time period. The report shows the workflow name, state, instance ID, and duration. You can view this report in the **Reports** view when you monitor a Data Integration Service or an application in the **Monitor** tab.

Minimum, Maximum, and Average Duration Report

Shows the minimum, maximum, and average duration for SQL data service and web service requests during a specified time period. You can view this report in the **Reports** view when you monitor a Data Integration Service, an SQL data service, a web service, or an application in the **Monitor** tab.

Most Active IP for SQL Data Service Requests

Shows the total number of SQL data service requests from each IP address during the specified time period. You can view this report in the **Reports** view when you monitor a Data Integration Service, an SQL data service, or an application in the **Monitor** tab.

Most Active SQL Data Service Connections

Shows SQL data service connections that received the most connection requests during the specified time period. The report shows the connection ID, SQL data service, and the total number of connection requests. You can view this report in the **Reports** view when you monitor a Data Integration Service, an application, or an SQL data service in the **Monitor** tab.

Most Active Users for Ad Hoc Jobs

Shows users that ran the most number of ad hoc jobs during the specified time period. The report shows the user name and the total number of jobs that the user ran. You can view this report in the **Reports** view when you monitor a Data Integration Service in the **Monitor** tab.

Most Active Web Service Client IP

Shows IP addresses that received the most number of web service requests during the specified time period. You can view this report in the **Reports** view when you monitor a Data Integration Service, an application, a web service, or web service operation in the **Monitor** tab.

Most Frequent Errors for Ad Hoc Jobs

Shows the most frequent errors for ad hoc jobs, regardless of job type, during the specified time period. The report shows the job type, error ID, and error count. You can view this report in the **Reports** view when you monitor a Data Integration Service in the **Monitor** tab.

Most Frequent Errors for SQL Data Service Requests

Shows the most frequent errors for SQL data service requests during the specified time period. The report shows the error ID and error count. You can view this report in the **Reports** view when you monitor a Data Integration Service, an SQL data service, or an application in the **Monitor** tab.

Most Frequent Faults for Web Service Requests

Shows the most frequent faults for web service requests during the specified time period. The report shows the fault ID and fault count. You can view this report in the **Reports** view when you monitor a Data Integration Service, a web service, or an application in the **Monitor** tab.

Summary Statistics View

You can view throughput and resource information for mapping jobs in the **Execution Statistics** view.

When you select an ad hoc mapping job, deployed mapping job, or mapping in a workflow in the contents panel, the details panel displays the **Summary Statistics** view. The **Summary Statistics** view displays run-time statistics about the throughput and resource usage for the job. The run-time period begins when the Data Integration Service starts reading from the first row.

You can view the following throughput statistics for the job:

- Source. The name of the mapping source file.
- Target name. The name of the target file.
- Rows. Number of rows read for source and target.
- Average Rows/Sec. Average number of rows read per second for source and target.
- Bytes. Number of bytes read for source and target.
- Average Bytes/Sec. Average number of bytes read per second for source and target.
- First Row Accessed. The date and time when the Data Integration Service started reading the first row in the source file.
- Dropped rows. Number of source rows that the Data Integration Service did not read.
- Rejected Rows. Number of target rows that the Data Integration Service did not write to target.

You can view the following resource usage statistics for the job:

- Executing node. Node where the Data Integration Service that ran the job is running.
- Average CPU usage. The average amount of CPU that the Data Integration Service used to run the job.
- Average Memory usage. The average amount of memory that the Data Integration Service used to run the job.

Detailed Statistics View

You can view graphs of the throughput and resource information for mapping jobs that run longer than one minute.

When you select an ad hoc mapping job, deployed mapping job, or mapping in a workflow in the contents panel of the **Execution Statistics** view, the details panel displays the **Detailed Statistics** view. The **Detailed Statistics** view displays run-time statistics about the throughput and resources usage for the job. The run-time period begins when the Data Integration Service starts reading from the first row.

The **Detailed Statistics** view displays the following graphs:

Throughput graph

Plots the number of rows read and written against the run time of the job.

CPU Usage graph

Plots the percentage of the Data Integration Service CPU that was allocated to run the job against the run time of the job.

Memory Usage graph

Plots the amount of memory, in megabytes, that the Data Integration Service allocated to run the job against the run time of the job.

Logs Tab

The **Logs** tab shows logs.

On the **Logs** tab, you can view the following types of logs:

- Domain log. Domain log events are log events generated from the domain functions that the Service Manager performs.
- Service log. Service log events are log events generated by each application service.
- User Activity log. User Activity log events monitor user activity in the domain.

The **Logs** tab displays the following components for each type of log:

- Filter. Configure filter options for the logs.
- Log viewer. Displays log events based on the filter criteria.
- Reset filter. Reset the filter criteria.
- Copy rows. Copy the log text of the selected rows.
- **Actions** menu. Contains options to save, purge, and manage logs. It also contains filter options.

Reports Tab

The **Reports** tab shows domain reports.

On the **Reports** tab, you can run the following domain reports:

- License Management Report. Run a report to monitor the number of software options purchased for a license and the number of times a license exceeds usage limits. Run a report to monitor the usage of logical CPUs and PowerCenter Repository Services. You run the report for a license.
- Web Services Report. Run a report to analyze the performance of web services running on a Web Services Hub. You run the report for a time interval.

Security Tab

You administer Informatica security on the Security tab of the Administrator tool.

The Security tab has the following components:

- Search section. Search for users, groups, or roles by name.
- Navigator. The Navigator appears in the left pane and displays groups, users, and roles.
- Contents panel. The contents panel displays properties and options based on the object selected in the Navigator and the tab selected in the contents panel.
- Security Actions menu. Contains options to create or delete a group, user, or role. You can manage LDAP configurations and operating system profiles. You can also view users that have privileges for a service.

Using the Search Section

Use the Search section to search for users, groups, and roles by name. Search is not case sensitive.

1. In the Search section, select whether you want to search for users, groups, or roles.
2. Enter the name or partial name to search for.

You can include an asterisk (*) in a name to use a wildcard character in the search. For example, enter "ad*" to search for all objects starting with "ad". Enter "*ad" to search for all objects ending with "ad".

3. Click Go.

The Search Results section appears and displays a maximum of 100 objects. If your search returns more than 100 objects, narrow your search criteria to refine the search results.

4. Select an object in the Search Results section to display information about the object in the contents panel.

Using the Security Navigator

The Navigator appears in the contents panel of the Security tab. When you select an object in the Navigator, the contents panel displays information about the object.

The Navigator on the Security tab displays one of the following sections based on what you are viewing:

- Groups section. Select a group to view the properties of the group, the users assigned to the group, and the roles and privileges assigned to the group.

- Users section. Select a user to view the properties of the user, the groups the user belongs to, and the roles and privileges assigned to the user.
- Roles section. Select a role to view the properties of the role, the users and groups that have the role assigned to them, and the privileges assigned to the role.
- Operating Profiles section. Select an operating profile to view the properties of the operating system profile, and the permissions assigned to users and groups that use the operating system profile.
- LDAP Configuration section. Select a configuration to view the LDAP server connection details, the LDAP security domain that contains users and groups imported from the LDAP directory service, and the LDAP synchronization schedule.

The Navigator provides different ways to complete a task. You can use any of the following methods to manage groups, users, and roles:

- Click the **Actions** menu. Each section of the Navigator includes an Actions menu to manage groups, users, roles, operating system profiles, or LDAP configurations.
- Right-click an object. Right-click an object in the Navigator to display the options available in the Actions menu.
- Use keyboard shortcuts. Use keyboard shortcuts to move to different sections of the Navigator.

Groups

A group is a collection of users and groups that can have the same privileges, roles, and permissions.

The Groups section of the Navigator organizes groups into security domain folders. A security domain is a collection of user accounts and groups in an Informatica domain. Native authentication uses the Native security domain which contains the users and groups created and managed in the Administrator tool. LDAP authentication uses LDAP security domains which contain users and groups imported from the LDAP directory service.

When you select a security domain folder in the Groups section of the Navigator, the contents panel displays all groups belonging to the security domain.

When you select a group in the Navigator, the contents panel displays the following tabs:

- Overview. Displays general properties of the group and users assigned to the group.
- Privileges. Displays the privileges and roles assigned to the group for the domain and for application services in the domain.
- Permissions. Displays the level of access that users within the group have perform tasks on domain objects, including nodes, grids and application services . Also displays the level of access that users within the group have to perform tasks on connection objects and operating system profiles.

Users

A user with an account in the Informatica domain can log in to the following application clients:

- Informatica Administrator
- PowerCenter Client
- Informatica Developer
- Informatica Analyst
- Metadata Manager

The Users section of the Navigator organizes users into security domain folders. A security domain is a collection of user accounts and groups in an Informatica domain. Native authentication uses the Native

security domain which contains the users and groups created and managed in the Administrator tool. LDAP authentication uses LDAP security domains which contain users and groups imported from the LDAP directory service.

When you select a security domain folder in the Users section of the Navigator, the contents panel displays all users belonging to the security domain.

When you select a user in the Navigator, the contents panel displays the following tabs:

- Overview. Displays general properties of the user and all groups to which the user belongs.
- Privileges. Displays the privileges and roles assigned to the user for the domain and for application services in the domain.
- Permissions. Displays the level of access that the user has to perform tasks on domain objects, including nodes, grids and application services . Also displays the level of access that the user has to perform tasks on connection objects and operating system profiles.

Roles

A role is a collection of privileges that you assign to a user or group. Privileges determine the actions that users can perform. You assign a role to users and groups for the domain and for application services in the domain.

The Roles section of the Navigator organizes roles into the following folders:

- System-defined Roles. Contains roles that you cannot edit or delete. The Administrator role is a system-defined role.
- Custom Roles. Contains roles that you can create, edit, and delete. The Administrator tool includes some custom roles that you can edit and assign to users and groups.

When you select a folder in the Roles section of the Navigator, the contents panel displays all roles belonging to the folder.

When you select a role in the Navigator, the contents panel displays the following tabs:

- Overview. Displays general properties of the role and the users and groups that have the role assigned for the domain and application services.
- Privileges. Displays the privileges assigned to the role for the domain and application services.

Operating System Profiles

An operating system profile is a security mechanism that the Data Integration Service and the PowerCenter Integration Service use to run mappings, workflows, and profiling jobs.

The Operating System Profiles section of the Navigator lists the operating system profiles configured in the domain.

When you select an operating system profile in the Navigator, the contents panel displays the following tabs:

- Properties. Displays general properties of the operating system profile configured for the Data Integration Service, for the PowerCenter Integration Service, or for both application services.
- Permissions. Displays the permissions assigned to users and groups that use the operating system profile. Also indicates whether the operating system profile is the default profile assigned to a user or group.

LDAP Configuration

You can configure an Informatica domain to enable users and groups imported from one or more LDAP directory services to log in to Informatica nodes, services, and application clients.

The LDAP Configuration section of the Navigator lists the LDAP configurations the domain uses.

When you select an LDAP configuration, the following tabs appear under the LDAP Configuration tab:

- **Overview.** Lists the connection details for the LDAP server that contains the directory service from which you want to import users and groups.
- **Security Domains.** Lists the details for the LDAP security domain that contains users and groups imported from the LDAP directory service.
- **Schedule.** Lists the details for the synchronization schedule specifying when the Service Manager updates the security domain with the users and groups in the LDAP directory service.

Account Management

To improve security in the Informatica domain, you can enforce lockout of user and administrator accounts after a specified number of failed login attempts.

The Account Lockout Configuration section of the Account Management page displays whether account lockout is enabled for user accounts and administrator accounts. The section also indicates the maximum number of failed login attempts allowed.

The Locked Out Native Users section of the page lists locked out user accounts in the native security domain. You can unlock a user account in the native security domain.

The Locked Out LDAP Users section of the page lists locked out user accounts in an LDAP security domain. You can unlock a user account in the Informatica domain. However, the LDAP administrator must unlock the user account in the LDAP server. The user cannot log in to the Informatica domain until the LDAP administrator unlocks the user account.

Audit Reports

Audit reports provide information about users and groups in the Informatica domain, and about the privileges, roles, and permissions assigned to each user or group.

You select the audit report to generate from the Select Report Type menu. You can generate the following audit reports:

User Personal Information

Displays contact information and status details of user accounts in the domain. You can select the users or groups for which you want to generate the report.

User Group Association

Displays information about users and the groups to which they belong. You can select the users or groups for which you want to generate the report.

Privileges

Displays information about privileges assigned to the users and groups in the domain. You can select the users or groups for which you want to generate the report.

Roles

Displays information about the roles assigned to the users and groups in the domain. You can select the roles for which you want to generate the report.




Domain Object Permissions

Displays information about the domain objects for which users and groups have permission. You can select the users or groups for which you want to generate the report.

Service States

You can identify the state of an application service by the icon that displays in the Administrator tool.

The following table describes the icons that indicate service states:






State	Icon
Available	
Disabled	
Unavailable	

Process States

You can identify the state of a Data Integration Service process or a PowerCenter Integration Service process by the icon that displays in the Administrator tool.

The state icons also indicate the type of node on which the process runs. If the primary node has high availability, a yellow diamond is superimposed on the process state icon. If the process runs on a grid, a grid icon is superimposed on the process state icon.

The following table describes the icons that indicate process states:

State	Icon
Aborted	
Aborted (with high availability)	
Aborted (Grid)	
Disabled	
Disabled (with high availability)	

State	Icon
Disabled (Grid)	
Failed	
Failed (with high availability)	
Failed (Grid)	
Running	
Running (with high availability)	
Running (Grid)	
Standing by or Delayed	
Standing by or Delayed (with high availability)	
Standing by or Delayed (Grid)	
Starting	
Starting (with high availability)	
Starting (Grid)	
Stopped	
Stopped (with high availability)	
Stopped (Grid)	
Stopping	
Stopping (with high availability)	
Stopping (Grid)	

Job States

You can identify the state of a job by the icon that displays in the Administrator tool.

The following table describes the icons associated with each job state:

State	Icon
Aborted	
Canceled	
Completed	
Failed	
Queued or Pending	
Running	
Starting	
Stopped	
Stopping	
Terminated	
Unknown	

Informatica Administrator Accessibility Overview

You can use a screen reader and keyboard shortcuts to navigate and work with the Administrator tool interface.

To turn the JAWS Virtual PC cursor on and off, use the keyboard shortcut **Insert+Z**.

Note: To use the JAWS screen reader with the Administrator tool, you must use Internet Explorer 11.

Keyboard Shortcuts

You can use keyboard shortcuts to navigate and work with the Administrator tool interface.

You can add, edit, and change values in the Administrator tool. Keyboard focus in the Administrator tool is indicated by a blue border around the interface label. A dotted line appears around a selected object indicating that the object is in focus. Tooltips appear when the label item receives keyboard focus or on mouse-over. The navigation order of objects in the editor is from top to bottom and left to right.

You can perform the following tasks with keyboard shortcuts:

Navigate among elements and select an element

Press **Tab**.

Select the previous object

Press **Shift+Tab**.

Navigate among perspective tabs

Press the **Left** or **Right** arrow key.

Select or clear a check box or radio button

Press the **Space** bar.

Upload files using the File Upload button

Press the **Space** bar.

Navigate through records in a dialog box

Press the **Up** or **Down** arrow key.

Select and open a drop-down menu item with sub-menus

Press the **Space** arrow key. To go back to the main menu, press **Esc**.

Edit the value of grid content such as the Access and Revoke fields in the Assign Permission and Edit Direct Permissions dialog box

Press the **Space** bar.

Note: You must enter appropriate values for all the form elements marked with an asterisk (*).

Move focus from Update Frequency drop-down menu to Time Range check box in the Statistics and Reports list grid in the Report and Statistic Settings dialog box of the Monitor tab or Monitoring tool

Press **Esc**.

You cannot access the split bars in the Administrator tool and increase or decrease the size of the panels using the keyboard. You cannot select multiple items with the **Ctrl** key in the Audit Reports tab under Security.

CHAPTER 4

Using the Domain View

This chapter includes the following topics:

- [About the Domain View, 65](#)
- [Dependency Graph, 66](#)
- [Command History, 68](#)
- [History View, 68](#)

About the Domain View

The Domain view displays an overview of the status of the domain and the objects it contains. You can use the Domain view to review current and historical information about the domain.


Use the Domain view to perform the following tasks:

- View current status, resource usage, and details for the domain and objects in the domain.
- View dependencies among objects in the domain.
- Perform domain actions such as shutting down the domain, enabling and disabling services, and shutting down nodes.
- View recent service commands that users issued from the Administrator tool.
- View historical information about status, resource usage, and events in the domain.

Dependency Graph

The **Dependency** graph displays dependencies among services, nodes, and grids in the Informatica domain.

The following image shows the **Dependency** graph for a Model Repository Service:



You can use the **Dependency** graph to perform the following tasks:

- View dependencies among nodes, services, and grids.
- Shut down a node.
- Enable, disable, or recycle a service.
- Disable or recycle services that depend on other services.

When you view dependencies for an object, the **Dependency** graph displays the upstream and downstream dependencies. Upstream dependencies are objects on which the selected object depends. Downstream dependencies are objects that depend on the selected object.

When you enable, disable, or recycle services from the **Dependency** graph, the actions appear in the **Command History** panel.

Viewing Dependencies for Application Services, Nodes, and Grids

You can view dependencies among application services, nodes, and grids in the Informatica domain.

1. In the Administrator tool, click the **Manage** tab.
2. In the contents panel, click the **Actions** menu for a domain object, and then select **View Dependencies**.

The **Dependency** graph opens and displays the object and its direct dependencies.

The **Dependency** graph displays domain objects connected by blue and orange lines, as follows:

- Blue lines indicate service-to-node and service-to-grid dependencies.
- Dashed blue lines indicate backup node-to-service dependencies.
- Orange lines indicate service-to-service dependencies, such as a Data Integration Service-to-Content Management Service dependency, or a Model Repository Service-to-Data Integration Service dependency.

The following table describes the information that appears in the **Dependency** graph based on the object:

Domain Object	Upstream Dependencies	Downstream Dependencies
Node	N/A	Services that run on the node.
Node running in a grid	N/A	The node has the following downstream dependencies: <ul style="list-style-type: none"> - Grid in which the node runs. - Service process that runs on the grid. - Service processes that run on the node, but not in the grid.
Service	Node on which the service process runs.	Services that depend on the service.
Service running on a grid	The service has the following upstream dependencies: <ul style="list-style-type: none"> - Nodes on which the service process runs. - The grid on which the service processes run. 	Services that depend on the service.
Service running in HA mode	Primary and backup nodes on which the service processes can run.	Services that depend on the service.
Grid	Nodes assigned to the grid.	Services that are running on the grid.

- In the **Dependency** graph, you can optionally complete the following tasks:
 - Select **Show Legend** to view information about the icons and lines used in the graph.
 - Click and drag to view different parts of the graph.
 - Zoom in or zoom out of the graph.
 - To exit the **Dependency** graph, click **Close**.

Recycling or Disabling Downstream Services

You can recycle or disable downstream services in the **Dependency** graph.

Downstream services are services that depend on other services. For example, a Data Integration Service depends on a Model Repository Service. You recycle or disable downstream services using the Actions menu for the service on which they depend. When you disable downstream services, the service processes abort.

- In the Administrator tool, click the **Manage** tab.
- Click the **Actions** menu for a domain object, and then select **View Dependencies**.
The **Dependency** graph opens and displays the object and its direct dependencies.
- Click **Actions > Recycle Downstream Dependents** or **Actions > Disable Downstream Dependents**.
The Recycle Downstream Dependents or Disable Downstream Dependents window appears.
- Optionally, choose whether the action is **Planned** or **Unplanned**.
- Optionally, enter comments about the action.
- Select the services that you would like to recycle or disable.
- Click **Recycle Services** or **Disable Services**.

Command History

The **Command History** panel on the **Domain** view displays recent service lifecycle commands that users issued from the Administrator tool. Service lifecycle commands include enable, disable, and recycle.

To view the command history, click **Domain Actions > View Command History**.

You can view the following information about the commands in the **Command History** panel:

- Service Name. Name of the service for which the command was issued.
- Service Type.
- Command.
- Status. Can be Failed, Success, or Queued.
- Status Updated
- Comments. Comments that users entered when they recycled or disabled the service.
- Message. Log messages associated with the command.

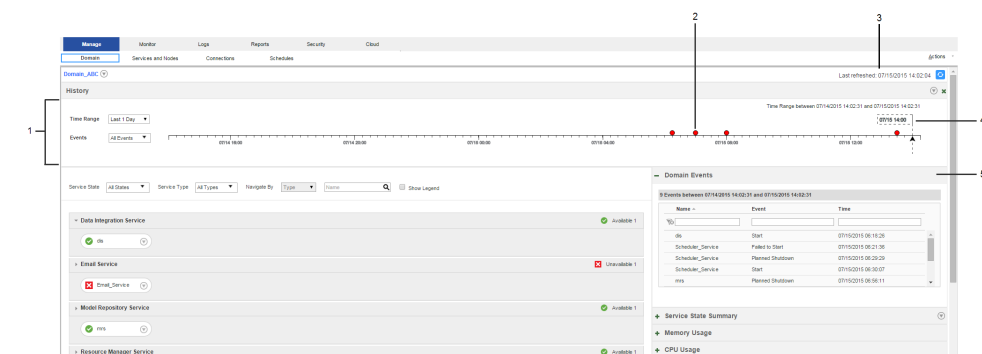
Optionally, you can show or hide columns in the command history. To change the columns, right-click the column header, and then select or clear columns.

Note: The command history is erased when you shut down or restart the master gateway node.

History View

The **History** view displays historical data for the domain, a service, or a node. You can view historical data for the contents panel, Service State Summary, resource usage indicators, and Details panel. You can also view information about events in the domain.

The following image shows the **History** view:



1. Time line
2. Event circle
3. Last Refreshed
4. Time line slider
5. Events panel

The **History** view has the following components:

Time line

Use the time line to choose which time range and events you would like to view. When events occur around a point in time, the point is marked with a red event circle. When you select a circle, the Events panel displays the events that occurred around that point in time. By default, the time line displays crashes and unplanned shutdowns that occurred in the last one day.

Events panel

Displays events that occurred during a specified time range. When you open the **History** view, the panel lists crashes and unplanned shutdowns that occurred in the last one day.

Contents panel

Displays current and historical domain contents and states. When you open the **History** view, the contents panel shows domain objects and states at the last refresh. When you drag the time line slider to a point in time, the contents panel shows domain objects and states at that point in time.

Service State Summary

Displays current and historical service state summaries. When you open the **History** view from the domain Actions menu, the Service State Summary appears and displays the number and states of services at the last refresh. When you drag the time line slider to a point in time, the Service State Summary displays the number of services that were available or unavailable at that time.

Details panel

Displays current and historical state for a service or node. When you open the **History** view, the Details panel displays the state of the service or node at the last refresh. When you drag the time line slider to a point in time, the Details panel displays the object state at that point in time.

Resource usage indicators

Display current and historical resource usage information. When you open the **History** view, the usage indicators display the usage statistics as they appeared at the last refresh. When you drag the time line slider to a point in time, the indicators display statistics for that point in time.

The data that you view in the History view is stored in the monitoring Model repository. Before you can view historical data, you must configure the monitoring Model repository in the **Manage > Services and Nodes > Monitoring Configuration** tab. You can view the per-minute data for up to two weeks in the past. After two weeks, the statistics that you view are hourly averages. You can view hourly averaged data for up to one year in the past.

Note: The monitoring Model Repository Service that you configure to store the historical data cannot capture data about events that occur while it is unavailable or disabled. For example, Start events for the monitoring Model Repository Service and Crash events for the node that the monitoring Model Repository Service process runs on do not appear in the History view.

Viewing History

You can view historical statistics for the domain, a service, or a node.

The amount of history that you can view depends on the monitoring Model repository options that you configure in the Monitoring Configuration tab. You must configure a monitoring Model repository in the Monitoring Configuration tab before you can view historical statistics.

1. To access the **History** view, click the **Actions** menu for the domain, a service, or a node, and then select **View History**.

The **History** view opens. The auto-refresh pauses at the Last Refreshed time stamp. The time line and **Events** panel display and display crashes and unplanned shutdowns that occurred in the last one day.

2. To change the time range, select a time range in the **Time Range** list.

When you change the time range, the time line refreshes and displays the selected time range.

3. To choose a custom time range, select **Custom** in the **Time Range** list.

You can select a custom time range from one hour to one month.

4. To change the events that display, select events in the **Events** list.

You can view information about the following events:

- Crash
- Failed to Start
- Unplanned Shutdown
- Planned Shutdown
- Start

When you select different events, the time line refreshes and reflects the change.

5. To view historical statistics about the domain, drag the time line slider to a point in time.

The **History** view refreshes and displays the domain status at that point in time.

6. To exit the **History** view, click **Close**.

Note: Current services display as Unavailable when the time range begins before they were created.

Viewing Events

You can view events for the domain, a node, or a service.

1. To view events, click the **Actions** menu for the domain, a node, or a service, and then choose **View History**.

The **History** view opens and displays the Events panel. By default, the Events panel lists crashes and unplanned shutdowns that have occurred in the last one day.

1. You can view the following information about events:

- Object name
- Object type
- Event type
- Time the event occurred
- Name of the associated node
- Comments about recycled or disabled services
- Comments about shut down nodes

2. Optionally, you can complete the following actions in the Events panel:

Option	Description
Search events	Type the search criteria in the search field, and then press enter. To clear the search, click the Reset Filters icon.
Sort a column	To sort a column in ascending order, click the column header. To sort the column in descending order, click the column header again.

Option	Description
Add or remove columns	To change the columns that appear in the Events panel, right-click a column header, and then select or deselect a column.
Reset filter	When you click a circle on the time line, you can clear the selection by clicking Reset filters.

CHAPTER 5

Domain Management

This chapter includes the following topics:

- [Domain Management Overview, 72](#)
- [Alert Management, 73](#)
- [Folder Management, 75](#)
- [Domain Security Management, 76](#)
- [User Security Management, 77](#)
- [Application Service Management, 77](#)
- [Gateway Configuration, 80](#)
- [Domain Configuration Management, 81](#)
- [Rename the Domain, 85](#)
- [Shutting Down a Domain, 85](#)
- [Domain Properties, 86](#)

Domain Management Overview

An Informatica domain is a collection of nodes and services that define the Informatica environment. To manage the domain, you manage the nodes and services within the domain.

Use the Administrator tool to complete the following tasks:

- Manage alerts. Configure, enable, and disable domain and service alerts for users.
- Create folders. Create folders to organize domain objects and manage security by setting permission on folders.
- Manage domain security. Configure secure communication between domain components.
- Manage user security. Assign privileges and permissions to users and groups.
- Manage application services. Enable, disable, recycle, and remove application services. Enable and disable service processes.
- Manage nodes. Configure node properties, such as the backup directory and resources, and shut down nodes.
- Configure gateway nodes. Configure nodes to serve as a gateway.
- Shut down the domain. Shut down the domain to complete administrative tasks on the domain.

- Manage domain configuration. Back up the domain configuration on a regular basis. You might need to restore the domain configuration from a backup to migrate the configuration to another database user account. You might also need to reset the database information for the domain configuration if it changes.
- Complete domain tasks. You can monitor the statuses of all application services and nodes, view dependencies among application services and nodes, and shut down the domain.
- Configure domain properties. For example, you can change the database properties, SMTP properties for alerts, and domain resiliency properties.

To manage nodes and services through a single interface, all nodes and services must be in the same domain. You cannot access multiple Informatica domains in the same Administrator tool window. You can share metadata between domains when you register or unregister a local repository in the local Informatica domain with a global repository in another Informatica domain.

Alert Management

Alerts provide users with domain and service alerts. Domain alerts provide notification about node failure and master gateway election. Service alerts provide notification about service process failover.

To use the alerts, complete the following tasks:

- Configure the SMTP settings for the outgoing email server.
- Subscribe to alerts.

After you configure the SMTP settings, users can subscribe to domain and service alerts.

Configuring SMTP Settings

You configure the SMTP settings for the outgoing mail server to enable alerts.

Configure SMTP settings on the domain **Properties** view.

1. In the Administrator tool, click the **Manage** tab.
2. Click the **Services and Nodes** view.
3. In the Navigator, select the domain.
4. In the contents panel, click the **Properties** view.
5. In the SMTP Configuration section, click **Edit**.
6. Edit the SMTP settings.

Property	Description
Host Name	The SMTP outbound mail server host name. For example, enter the Microsoft Exchange Server for Microsoft Outlook.
Port	Port used by the outgoing mail server. Valid values are from 1 to 65535. Default is 25.
User name	The user name for authentication upon sending if required by the outbound mail server.

Property	Description
Password	The user password for authentication upon sending if required by the outbound mail server.
Sender Email Address	The email address that the Service Manager uses in the From field when sending notification emails. If you leave this field blank, the Service Manager uses Administrator@<host name> as the sender.

7. Click **OK**.

Subscribing to Alerts

After you complete the SMTP configuration, you can subscribe to alerts.

1. Verify that the domain administrator has entered a valid email address for your user account on the **Security** page.
If the email address or the SMTP configuration is not valid, the Service Manager cannot deliver the alert notification.
2. In the Administrator tool header area, click **Manage > Preferences**.
The **Preferences** page appears.
3. In the User Preferences section, click **Edit**.
The **Edit Preferences** dialog box appears.
4. Select **Subscribe for Alerts**.
5. Click **OK**.
6. Click **OK**.

The Service Manager sends alert notification emails based on your domain privileges and permissions.

The following table lists the alert types and events for notification emails:

Alert Type	Event
Domain	Node Failure Master Gateway Election
Service	Service Process Failover

Viewing Alerts

When you subscribe to alerts, you can receive domain and service notification emails for certain events. When a domain or service event occurs that triggers a notification, you can track the alert status in the following ways:

- The Service Manager sends an alert notification email to all subscribers with the appropriate privilege and permission on the domain or service.
- The Log Manager logs alert notification delivery success or failure in the domain or service log.

For example, the Service Manager sends the following notification email to all alert subscribers with the appropriate privilege and permission on the service that failed:

```
From: Administrator@<database host>
To: Jon Smith
```

```
Subject: Alert message of type [Service] for object [HR_811].
The service process on node [node01] for service [HR_811] terminated unexpectedly.
```

In addition, the Log Manager writes the following message to the service log:

```
ALERT_10009 Alert message [service process failover] of type [service] for object
[HR_811] was successfully sent.
```

You can review the domain or service logs for undeliverable alert notification emails. In the domain log, filter by Alerts as the category. In the service logs, search on the message code ALERT. When the Service Manager cannot send an alert notification email, the following message appears in the related domain or service log:

```
ALERT_10004: Unable to send alert of type [alert type] for object [object name], alert
message [alert message], with error [error].
```

Folder Management

Use folders in the domain to organize objects and to manage security.

Folders can contain nodes, services, grids, licenses, and other folders. You might want to use folders to group services by type. For example, you can create a folder called IntegrationServices and move all Integration Services to the folder. Or, you might want to create folders to group all services for a functional area, such as Sales or Finance.

When you assign a user permission on the folder, the user inherits permission on all objects in the folder.

You can perform the following tasks with folders:

- View services and nodes. View all services in the folder and the nodes where they run. Click a node or service name to access the properties for that node or service.
- Create folders. Create folders to group objects in the domain.
- Move objects to folders. When you move an object to a folder, folder users inherit permission on the object in the folder. When you move a folder to another folder, the other folder becomes a parent of the moved folder.
- Remove folders. When you remove a folder, you can delete the objects in the folder or move them to the parent folder.

Note: The System_Services folder is created for you when you create the domain, and contains all of the system services. A system service is an application service that can have a single instance in the domain. You cannot delete, move, or edit the properties or contents of the System_Services folder.

Creating a Folder

You can create a folder in the domain or in another folder.

1. In the Administrator tool, click the **Manage** tab > **Services and Nodes** view.
2. In the Domain Navigator, select the domain or folder in which you want to create a folder.
3. On the Navigator Actions menu, click New > Folder.

4. Edit the following properties:

Node Property	Description
Name	Name of the folder. The name is not case sensitive and must be unique within the domain. It cannot exceed 80 characters or begin with @. It also cannot contain spaces or the following special characters: ` ~ % ^ * + = { } \ ; : ' " / ? . , < > ! () [
Description	Description of the folder. The description cannot exceed 765 characters.
Path	Location in the Navigator.

5. Click OK.

Moving Objects to a Folder

When you move an object to a folder, folder users inherit permission on the object. When you move a folder to another folder, the moved folder becomes a child object of the folder where it resides.

Note: The domain serves as a folder when you move objects in and out of folders.

1. In the Administrator tool, click the **Manage** tab > **Services and Nodes** view.
2. In the Domain Navigator, select an object.
3. On the Navigator Actions menu, select Move to Folder.
4. In the Select Folder dialog box, select a folder, and click OK.

Removing a Folder

When you remove a folder, you can delete the objects in the folder or move them to the parent folder.

1. In the Administrator tool, click the **Manage** tab > **Services and Nodes** view.
2. In the Domain Navigator, select a folder.
3. On the Navigator Actions menu, select Delete.
4. Confirm that you want to delete the folder.

You can delete the contents only if you have the appropriate privileges and permissions on all objects in the folder.

5. Choose to wait until all processes complete or to abort all processes.
6. Click OK.

Domain Security Management

You can configure Informatica domain components to use the Secure Sockets Layer (SSL) protocol or the Transport Layer Security (TLS) protocol to encrypt connections with other components. When you enable SSL or TLS for domain components, you ensure secure communication.

You can configure secure communication in the following ways:

Between services within the domain

You can configure secure communication between services within the domain.

Between the domain and external components

You can configure secure communication between Informatica domain components and web browsers or web service clients.

Each method of configuring secure communication is independent of the other methods. When you configure secure communication for one set of components, you do not need to configure secure communication for any other set.

Note: If you change a secure domain to a non-secure domain or from a non-secure domain to a secure domain, you must delete the domain configuration in the Developer tool and PowerCenter client tools and configure the domain again in the client.

User Security Management

You manage user security within the domain with privileges and permissions.

Privileges determine the actions that users can complete on domain objects. Permissions define the level of access a user has to a domain object. Domain objects include the domain, folders, nodes, grids, licenses, database connections, operating system profiles, and application services.

Even if a user has the domain privilege to complete certain actions, the user might also require permission to complete the action on a particular object. For example, a user has the Manage Services domain privilege which grants the user the ability to edit application services. However, the user also must have permission on the application service. A user with the Manage Services domain privilege and permission on the Development Repository Service but not on the Production Repository Service can edit the Development Repository Service but not the Production Repository Service.

To log in to the Administrator tool, a user must have the Access Informatica Administrator domain privilege. If a user has the Access Informatica Administrator privilege and permission on an object, but does not have the domain privilege that grants the ability to modify the object type, then the user can view the object. For example, if a user has permission on a node, but does not have the Manage Nodes and Grids privilege, the user can view the node properties but cannot configure, shut down, or remove the node.

If a user does not have permission on a selected object in the Navigator, the contents panel displays a message indicating that permission on the object is denied.

Application Service Management

You can perform the following common administration tasks for application services:

- Enable and disable services and service processes.
- Configure the domain to restart service processes.
- Remove an application service.
- Troubleshoot problems with an application service.

Note: You can perform all of the common administration tasks for system services, except for removing the system service.

Enabling and Disabling Services and Service Processes

You can enable and disable application services and service processes in the Administrator tool. When a service is enabled, there must be at least one service process enabled and running for the service to be available. By default, all service processes are enabled.

The behavior of a service when it starts service processes depends on how it is configured:

- If the service is configured for high availability, then the service starts the service process on the primary node. The service processes on the backup nodes are in Standing By state.
- If the service is configured to run on a grid, then the service starts service processes on all nodes that have the service role.

A service does not start a disabled service process in any situation.

The state of a service depends on the state of its processes. A service can have the following states:

- Available. You have enabled the service and at least one service process is running. The service is available to process requests.
- Unavailable. You have enabled the service and none of its processes are running. This can be because the service processes are disabled or failed to start. The service is not available to process requests.
- Disabled. You have disabled the service.

You can disable a service to perform a management task, such as changing the data movement mode for a PowerCenter Integration Service. You might want to disable the service process on a node if you need to shut down the node for maintenance. When you disable a service, all associated service processes stop, but they remain enabled.

The following table describes the different states of a service and its processes:

Service Process Configuration	Service Process State	Description
Enabled	Running	The service process is running on the node.
Enabled	Standing By	The service process is enabled but is not running because another service process is running as the primary service process. It is on standby to run in case of service failover.
Disabled	Disabled	The service is enabled but the service process is not running on the node.
Enabled	Stopped	The service is unavailable.
Enabled	Failed	The service and service process are enabled, but the service process could not start.

Viewing Service Processes

You can view the state of a service process on the Processes view of a service. You can view the state of all service processes on the Overview view of the domain.

To view the state of a service process:

1. In the Administrator tool, click the **Manage** tab > **Services and Nodes** view.
2. In the Domain Navigator, select a service.
3. In the contents panel, select the Processes view.

The Processes view displays the state of the processes.

Configuring Restart for Service Processes

If an application service process becomes unavailable while a node is running, the domain tries to restart the process on the same node based on the restart options configured in the domain properties.

1. In the Administrator tool, click the **Manage** tab > **Services and Nodes** view.
2. In the Domain Navigator, select the domain.
3. In the Properties view, configure the following restart properties:

Domain Property	Description
Maximum Restart Attempts	Number of times within a specified period that the domain attempts to restart an application service process when it fails. The value must be greater than or equal to 1. Default is 3.
Within Restart Period (sec)	Maximum period of time that the domain spends attempting to restart an application service process when it fails. If a service fails to start after the specified number of attempts within this period of time, the service does not restart. Default is 900.

Removing Application Services

You can remove an application service using the Administrator tool. Before removing an application service, you must disable it.

Note: You cannot remove a system service.

Disable the service before you delete the service to ensure that the service is not running any processes. If you do not disable the service, you may have to choose to wait until all processes complete or abort all processes when you delete the service.

1. In the Administrator tool, click the **Manage** tab > **Services and Nodes** view.
2. In the Domain Navigator, select the application service.
3. In the **Manage** tab Actions menu, select **Delete**.
4. In the warning message that appears, click **Yes** to stop other services that depend on the application service.
5. If the **Disable Service** dialog box appears, choose to wait until all processes complete or abort all processes, and then click **OK**.

Troubleshooting Application Services

I think that a service is using incorrect environment variable values. How can I find out which environment variable values are used by a service.

Set the error severity level for the node to debug. When the service starts on the node, the Domain log will display the environment variables that the service is using.

Gateway Configuration

A domain requires at least one node configured as a gateway node. You can configure multiple gateway nodes as backups.

One gateway node in the domain serves as the master gateway node for the domain. The Service Manager on the master gateway node accepts service requests and manages the domain and services in the domain. If the domain has one gateway node and it becomes unavailable, the domain cannot accept service requests. If the domain has multiple gateway nodes and the master gateway node becomes unavailable, the Service Managers on the other gateway nodes elect a new master gateway node. The new master gateway node accepts service requests. Only one gateway node can be the master gateway node at any particular time.

You can make the following changes to the gateway configuration for the domain:

Convert a worker node to serve as a gateway node.

You can convert a worker node to serve as a gateway node if the worker node is running and has the service role enabled. When you convert a worker node to a gateway node, you must specify the log directory for the node. If you have multiple gateway nodes, configure all gateway nodes to write log files to the same directory on a shared disk.

After you convert a worker node to a gateway node, the Service Manager on the master gateway node writes the domain configuration database connection to the nodemeta.xml file of the new gateway node.

Convert a gateway node to serve as a worker node.

You can convert a gateway node to serve as a worker node if another node in the domain is configured as a gateway node.

If you convert a master gateway node to serve as a worker node, you must restart the node to make the Service Managers elect a new master gateway node. If you do not restart the node, the node continues as the master gateway node until you restart the node or the node becomes unavailable.

Configuring the Gateway and Worker Nodes

You can convert an existing worker node to a gateway node. Or, you can convert an existing gateway node to a worker node.

1. In the Administrator tool, click the **Manage** tab > **Services and Nodes** view.
2. In the Domain Navigator, select the domain.
3. In the contents panel, select the **Properties** view.
4. In the **Properties** view, click **Edit** in the **Gateway Configuration Properties** section.
5. To convert a worker node to a gateway node, complete the following steps:
 - a. Select the check box next to the node.
 - b. If the domain uses a secure domain configuration database, specify the truststore file and password for the database.
 - c. Configure the directory path for the log files for each node that you convert to a gateway node.
If you have multiple gateway nodes, configure all gateway nodes to write log files to the same directory on a shared disk.

Note: You must use the `infacmd isp SwitchToWorkerNode` command to convert a worker node to a gateway node in a domain configured to use SAML authentication. See the *Informatica Command Reference* for instructions on using the `infacmd isp SwitchToWorkerNode` command.

6. To convert a gateway node to a worker node, clear the check box next to the node.

7. Click **OK**.

Domain Configuration Management

The Service Manager on the master gateway node manages the domain configuration. The domain configuration is a set of metadata tables stored in a relational database that is accessible by all gateway nodes in the domain. Each time you make a change to the domain, the Service Manager writes the change to the domain configuration. For example, when you add a node to the domain, the Service Manager adds the node information to the domain configuration. The gateway nodes use a JDBC connection to access the domain configuration database.

You can perform the following domain configuration management tasks:

- Back up the domain configuration. Back up the domain configuration on a regular basis. You may need to restore the domain configuration from a backup if the domain configuration in the database becomes corrupt.
- Restore the domain configuration. You may need to restore the domain configuration if you migrate the domain configuration to another database user account. Or, you may need to restore the backup domain configuration to a database user account.
- Migrate the domain configuration. You may need to migrate the domain configuration to another database user account.
- Configure the connection to the domain configuration database. Each gateway node must have access to the domain configuration database. You configure the database connection when you create a domain. If you change the database connection information or migrate the domain configuration to a new database, you must update the database connection information for each gateway node.
- Configure custom properties. Configure domain properties that are unique to your environment or that apply in special cases. Use custom properties only if Informatica Global Customer Support instructs you to do so.

Note: The domain configuration database and the Model repository cannot use the same database user schema.

Backing Up the Domain Configuration

Back up the domain configuration on a regular basis. You may need to restore the domain configuration from a backup file if the domain configuration in the database becomes corrupt.

Run the *infasetup BackupDomain* command to back up the domain configuration to a binary file.

Note: If the *infasetup BackupDomain* command fails with a Java memory error, increase the system memory available for *infasetup*. To increase system memory, set the `-Xmx` value in the `INFA_JAVA_COMD_OPTS` environment variable.

When you run this command, *infasetup* backs up the domain configuration database tables. To restore the domain to another database, you must back up the `ISP_RUN_LOG` table contents manually to get the previous workflow and session logs.

Additionally, use the database backup utility to manually back up additional repository tables that the *infasetup* command does not back up.

Restoring the Domain Configuration

You can restore domain configuration from a backup file. You may need to restore the domain configuration if the domain configuration in the database becomes inconsistent or if you want to migrate the domain configuration to another database.

Informatica restores the domain configuration from the current version. If you have a backup file from an earlier product version, you must use the earlier version to restore the domain configuration.

You can restore the domain configuration to the same or a different database user account. If you restore the domain configuration to a database user account with existing domain configuration, you must configure the command to overwrite the existing domain configuration. If you do not configure the command to overwrite the existing domain configuration, the command fails.

Each node in a domain has a host name and port number. When you restore the domain configuration, you can disassociate the host names and port numbers for all nodes in the domain. You might do this if you want to run the nodes on different machines. After you restore the domain configuration, you can assign new host names and port numbers to the nodes. Run the *infasetup* DefineGatewayNode or DefineWorkerNode command to assign a new host name and port number to a node.

If you restore the domain configuration to another database, you must reset the database connections for all gateway nodes.

Important: You lose all data in the summary tables when you restore the domain configuration.

Complete the following tasks to restore the domain:

1. Disable the application services. Disable the application services in complete mode to ensure that you do not abort any running service process. You must disable the application services to ensure that no service process is running when you shut down the domain.
2. Shut down the domain. You must shut down the domain to ensure that no change to the domain occurs while you are restoring the domain.
3. Run the *infasetup* RestoreDomain command to restore the domain configuration to a database. The RestoreDomain command restores the domain configuration in the backup file to the specified database user account.
4. Assign new host names and port numbers to the nodes in the domain if you disassociated the previous host names and port numbers when you restored the domain configuration. Run the *infasetup* DefineGatewayNode or DefineWorkerNode command to assign a new host name and port number to a node.
5. Reset the database connections for all gateway nodes if you restored the domain configuration to another database. All gateway nodes must have a valid connection to the domain configuration database.

Migrating the Domain Configuration

You can migrate the domain configuration to another database user account. You may need to migrate the domain configuration if you no longer support the existing database user account. For example, if your company requires all departments to migrate to a new database type, you must migrate the domain configuration.

1. Shut down all application services in the domain.
2. Shut down the domain.
3. Back up the domain configuration.
4. Create the database user account where you want to restore the domain configuration.

5. Restore the domain configuration backup to the database user account.
6. Update the database connection for each gateway node.
7. Start all nodes in the domain.
8. Enable all application services in the domain.

Important: Summary tables are lost when you restore the domain configuration.

Step 1. Disable All Application Services

You must disable all application services to disable all service processes. If you do not disable an application service and a user starts a service process while you are backing up and restoring the domain, the service process changes may be lost and data may become corrupt.

Disable application services in complete mode to ensure that you do not abort running service processes.

Disable the application services in the following order:

1. Web Services Hub
2. SAP BW Service
3. Metadata Manager Service
4. PowerCenter Integration Service
5. PowerCenter Repository Service
6. Search Service
7. Analyst Service
8. Content Management Service
9. Data Integration Service
10. Model Repository Service

Step 2. Shut Down the Domain

You must shut down the domain to ensure that users do not modify the domain while you are migrating the domain configuration. For example, if the domain is running when you are backing up the domain configuration, users can create new services and objects. Also, if you do not shut down the domain and you restore the domain configuration to a different database, the domain becomes inoperative. The connections between the gateway nodes and the domain configuration database become invalid. The gateway nodes shut down because they cannot connect to the domain configuration database. A domain is inoperative if it has no running gateway node.

Step 3. Back Up the Domain Configuration

Run the `infasetup BackupDomain` command to back up the domain configuration to a binary file.

Step 4. Create a Database User Account

Create a database user account if you want to restore the domain configuration to a new database user account.

Step 5. Restore the Domain Configuration

Run the *infasetup* RestoreDomain command to restore the domain configuration to a database. The RestoreDomain command restores the domain configuration in the backup file to the specified database user account.

Step 6. Update the Database Connection

If you restore the domain configuration to a different database user account, you must update the database connection information for each gateway node in the domain. Gateway nodes must have a connection to the domain configuration database to retrieve and update domain configuration.

Step 7. Start All Nodes in the Domain

Start all nodes in the domain. You must start the nodes to enable services to run.

1. Shut down the gateway node that you want to update.
2. Run the *infasetup* UpdateGatewayNode command to update the gateway node.
3. Start the gateway node.
4. Repeat this process for each gateway node.

Step 8. Enable All Application Services

Enable all application services that you previously shut down. Application services must be enabled to run service processes.

Updating the Domain Configuration Database Connection

All gateway nodes must have a connection to the domain configuration database to retrieve and update domain configuration. When you create a gateway node or configure a node to serve as a gateway, you specify the database connection, including the database user name and password. If you migrate the domain configuration to a different database or change the database user name or password, you must update the database connection for each gateway node. For example, as part of a security policy, your company may require you to change the password for the domain configuration database every three months.

To update the node with the new database connection information, complete the following steps:

1. Shut down the gateway node.
2. Run the *infasetup* UpdateGatewayNode command.

If you change the user or password, you must update the node.

To update the node after you change the user or password, complete the following steps:

1. Shut down the gateway node.
2. Run the *infasetup* UpdateGatewayNode command.

If you change the host name or port number, you must redefine the node.

To redefine the node after you change the host name or port number, complete the following steps:

1. Shut down the gateway node.
2. In the Administrator tool, remove the node association.
3. Run the *infasetup* DefineGatewayNode command.

Rename the Domain

You can change the domain name and update nodes to reference the updated domain name.

If the Informatica domain uses Kerberos authentication, all service and node SPNs have the same Kerberos realm name. After you change the Informatica domain name, you must generate SPNs and keytab files with the new Informatica domain name.

To rename the domain, complete the following tasks:

1. If the domain contains a PowerCenter global repository, you must unregister all local repositories from the global repository.
2. Shut down the domain. Shut down the domain through the Administrator tool, ensuring that all nodes are shut down.
3. Back up the domain with the `infasetup BackupDomain` command.
4. Back up the sitekey and keytab files. By default, the files are in the following location:

```
<Informatica installation directory>/isp/config/keys
```
5. Update the domain and nodes.
To update the domain name, run the `infasetup updateDomainName` command from any gateway node.
Run the `updateGatewayNode` and `updateWorkerNode` commands with the updated domain name for all the gateway and worker nodes.
6. On PowerCenter, register the local repository with a connected global repository with the updated domain name with the `pmrep Register` command.
7. You can create SPN and keytab files with the updated domain name for Kerberos authentication. Copy the keytab files in the keys directory. You can continue to use the older site key file. If you need to regenerate the site key when it is missing or corrupted, you must provide the older domain name.
8. Optionally, you can run the License Management Report in the Administrator tool to review the updated domain name.
9. You must configure the Informatica clients to use the updated domain name.

Shutting Down a Domain

To run administrative tasks on a domain, you might need to shut down the domain. For example, to back up and restore a domain configuration, you must first shut down the domain.

When you shut down a domain, the Service Manager on the master gateway node stops all application services and Informatica services in the domain. Any service processes running on nodes in the domain are aborted. To avoid possible loss of data or metadata and allow running processes to complete, you can shut down each node from the Administrator tool or from the operating system.

Before you shut down a domain, verify that all processes, including workflows, have completed and no users are logged in to repositories in the domain.

1. Click the **Manage** tab.
2. Click the **Services and Nodes** view.
3. In the Navigator, select the domain.
4. Click **Manage tab Actions > Shut Down Domain** .

The **Shut Down Domain** dialog box lists the processes that are running in the domain.

5. Click **Shut down**.
The **Shutdown Down Domain** dialog box displays a warning message.
6. Click **Shut down**.
The Service Manager on the master gateway node shuts down the application services and Informatica services on each node in the domain.
7. To restart the domain, restart Informatica services on the gateway and worker nodes in the domain.

Domain Properties

On the **Manage** tab, you can configure domain properties including database properties, gateway configuration, and service levels.

To view and edit properties, click the **Manage** tab. In the Navigator, select a domain. Then click the **Properties** view in the contents panel. The contents panel shows the properties for the domain.

You can configure the properties to change the domain. You cannot change the database properties in the Admin Console. You need to change these properties using the Command `UpdateGatewayNode`. You can change SMTP properties for alerts, and the domain resiliency properties.

You can configure the following domain properties:

- General properties. Edit general properties, such as service resilience and dispatch mode.
- Database properties. View the database properties, such as database name and database host.
- Gateway configuration properties. Configure a node to serve as gateway and specify the location to write log events.
- Service level management. Create and configure service levels.
- SMTP configuration. Edit the SMTP settings for the outgoing mail server to enable alerts.
- Custom properties. Edit custom properties that are unique to the Informatica environment or that apply in special cases. When you create a domain, it has no custom properties. Use custom properties only at the request of Informatica Global Customer Support.

General Properties

In the General Properties area, you can configure general properties for the domain.

To edit general properties, click **Edit**.

The following table describes the properties that you can edit in the General Properties area:

Property	Description
Name	Read-only. The name of the domain.
Resilience Timeout	The number of seconds that an application service tries to connect or reconnect to the PowerCenter Repository Service or the PowerCenter Integration Service. Valid values are from 0 to 1000000. Default is 30 seconds.

Property	Description
Limit on Resilience Timeout	The maximum number of seconds that application clients or application services can try to connect or reconnect to the PowerCenter Repository Service or the PowerCenter Integration Service. Default is 180 seconds.
Restart Period	The maximum amount of time in seconds that the domain spends trying to restart an application service process. Valid values are from 0 to 1000000.
Maximum Restart Attempts within Restart Period	The number of times that the domain tries to restart an application service process. Valid values are from 0 to 1000. If you set the value as 0, the domain does not try to restart the service process.
Dispatch Mode	The mode that the Load Balancer uses to dispatch PowerCenter Integration Service tasks to nodes in a grid. Select one of the following dispatch modes: <ul style="list-style-type: none"> - MetricBased - RoundRobin - Adaptive
Enable Secure Communication	Configures services to use the TLS protocol to transfer data securely within the domain. When you enable secure communication for the domain, services use secure connections to communicate with other Informatica application services and clients. Verify that all domain nodes are available before you enable secure communication for the domain. If a node is not available, the secure communication changes cannot be applied to the Service Manager of the node. To apply changes, restart the domain. Set this property to True or False.
Service Resilience Timeout	The maximum number of seconds that application clients and application services can try to connect to the Data Integration Service or to the Model Repository Service. The default is 180 seconds.

Database Properties

In the Database Properties area, you can view the database properties for the domain, such as database name and database host. You cannot edit these properties in the Admin Console. You need to update these properties using the command `UpdateGatewayNode`.

The following table describes the Database properties :

Property	Description
Database Type	The type of database that stores the domain configuration metadata.
Database Host	The name of the machine hosting the database.
Database Port	The port number used by the database.
Database Name	The name of the database.
Database User	The user account for the database containing the domain configuration information.
Database TLS enabled	Indicates whether the database for the domain configuration repository is a secure database. True if the domain configuration repository database is secure. You can use a secure domain configuration repository if secure communication is enabled for the Informatica domain.

Note: The service manager uses the DataDirect drivers included with the Informatica installation. Informatica does not support the use of any other database driver.

Gateway Configuration Properties

In the Gateway Configuration Properties area, you can configure a node to serve as gateway for a domain and specify the directory where the Service Manager on this node writes the log event files.

If you edit gateway configuration properties, previous logs do not appear. Also, the changed properties apply to restart and failover scenarios only.

To edit gateway configuration properties, click **Edit**.

To sort gateway configuration properties, click the header of the column by which you want to sort.

The following table describes the properties that you can edit in the Gateway Configuration Properties area:

Property	Description
Node Name	Read-only. The name of the node.
Status	The status of the node.
Gateway	To configure the node as a gateway node, select this option. If the domain uses a secure domain configuration database, you must specify the truststore file and password for the database. To configure the node as a worker node, clear this option.
Log Directory Path	The directory path for the log event files. If the Log Manager cannot write to the directory path, it writes log events to the node.log file on the master gateway node.

Secure Domain Configuration Repository

If you configure a node as a gateway node and the domain uses a secure domain configuration database, you must specify the truststore file and password for the secure database.

If you configure multiple gateway nodes for the domain, set the database truststore file and password for all gateway nodes.

The following table describes the database truststore properties:

Property	Description
Database Truststore Password	Password for the truststore file.
Database Truststore Location	Path and file name of the truststore file for the secure database.

Note: To use a secure domain configuration repository database, the secure communication option must be enabled for the domain.

Service Level Management

In the Service Level Management area, you can view, add, and edit service levels.

Service levels set priorities among tasks that are waiting to be dispatched. When the Load Balancer has more tasks to dispatch than the PowerCenter Integration Service can run at the time, the Load Balancer places

those tasks in the dispatch queue. When multiple tasks are in the dispatch queue, the Load Balancer uses service levels to determine the order in which to dispatch tasks from the queue.

Because service levels are domain properties, you can use the same service levels for all repositories in a domain. You create and edit service levels in the domain properties or by using infacmd.

You can edit but you cannot delete the Default service level, which has a dispatch priority of 5 and a maximum dispatch wait time of 1800 seconds.

To add a service level, click **Add**.

To edit a service level, click the link for the service level.

To delete a service level, select the service level and click the Delete button.

The following table describes the properties that you can edit in the Service Level Management area:

Property	Description
Name	The name of the service level. The name is not case sensitive and must be unique within the domain. It cannot exceed 128 characters or begin with the @ character. It also cannot contain spaces or the following special characters: <code>` ~ % ^ * + = { } \ ; : / ? . < > ! ()] [</code> After you add a service level, you cannot change its name.
Dispatch Priority	A number that sets the dispatch priority for the service level. The Load Balancer dispatches high priority tasks before low priority tasks. Dispatch priority 1 is the highest priority. Valid values are from 1 to 10. Default is 5.
Maximum Dispatch Wait Time (seconds)	The amount of time in seconds that the Load Balancer waits before it changes the dispatch priority for a task to the highest priority. Setting this property ensures that no task waits forever in the dispatch queue. Valid values are from 1 to 86400. Default is 1800.

SMTP Configuration

Use the SMTP Configuration properties to configure SMTP settings for the domain. The outgoing mail server uses the SMTP settings to send alerts and scorecard notifications.

The following table describes the properties that you can edit in the SMTP Configuration area:

Property	Description
Host Name	The SMTP outbound mail server host name. For example, enter the Microsoft Exchange Server for Microsoft Outlook.
Port	Port used by the outgoing mail server. Valid values are from 1 to 65535. Default is 25.
User name	The user name for authentication upon sending if required by the outbound mail server.
Password	The user password for authentication upon sending if required by the outbound mail server.
Sender Email Address	The email address that the Service Manager uses in the From field when sending notification emails. If you leave this field blank, the Service Manager uses Administrator@<host name> as the sender.

Custom Properties for the Domain

Configure custom properties that are unique to specific environments.

You might need to apply custom properties in special cases. When you define a custom property, enter the property name and an initial value. Define custom properties only at the request of Informatica Global Customer Support.

CHAPTER 6

Nodes

This chapter includes the following topics:

- [Nodes Overview, 91](#)
- [Node Types, 92](#)
- [Node Roles, 93](#)
- [Define and Add Nodes, 95](#)
- [Configuring Node Properties, 96](#)
- [Shutting Down and Restarting the Node, 98](#)
- [Removing the Node Association, 100](#)
- [Removing a Node, 100](#)

Nodes Overview

A node is the logical representation of a machine in the domain. When you configure a domain with multiple nodes, you can scale service processing across multiple machines. The Service Manager runs on all nodes in the domain to support the domain and application services. If the Service Manager is not running, the node is not available.

An installation on multiple machines consists of a master gateway node, which hosts the domain, and additional gateway nodes and worker nodes that run Informatica application services. The node type determines whether the node can serve as a gateway or a worker node and determines the domain functions that the node performs. You define the node type when you install Informatica services and join the node to the domain. You can use the Administrator tool to change the node type after installation.

By default, each node in the domain can run application services and computational processes. The node role determines whether a node can run application services, computational processes, or both. If the node has the service role, you can view the application service processes running on the node. Before you remove or shut down a node, verify that all running processes are stopped. You might need to shut down the node if you need to perform maintenance on the machine or to ensure that domain configuration changes take effect.

Use the Manage tab of the Administrator tool to manage nodes, including configuring node properties, updating a node role, and removing nodes from a domain. The properties that you can configure depend on the node role.

If your license includes grid, you can configure the Data Integration Service or the PowerCenter Integration Service to run on a grid. A grid is an alias assigned to a group of nodes. When you run jobs on a grid of nodes, you improve scalability and performance by distributing jobs to processes running on multiple nodes in the grid. When the PowerCenter Integration Service runs on a grid, you can configure it to check the resources

available on each node. Assign connection resources and define custom and file/directory resources on a node that is assigned to a PowerCenter Integration Service grid.

Node Types

The node type determines whether the node can serve as a gateway or worker node and determines the domain functions that the node performs.

You define the node type when you install Informatica services and join the node to the domain. You can use the Administrator tool to change the node type after installation. You change the node type in the gateway configuration properties for the domain.

RELATED TOPICS:

- [“Gateway Configuration” on page 80](#)

Gateway Nodes

A gateway node is any node that you configure to serve as a gateway for the domain. A gateway node can run application services and perform computations, and it can serve as a master gateway node. One gateway node acts as the master gateway at any given time. The master gateway node is the entry point to the domain.

The Service Manager on the master gateway node performs all domain functions on the master gateway node. The Service Managers running on other gateway nodes perform limited domain functions on those nodes.

You can configure more than one node to serve as a gateway. If the master gateway node becomes unavailable, the Service Managers on other gateway nodes elect another master gateway node. If you configure only one node to serve as the gateway and the node becomes unavailable, the domain cannot accept service requests.

Worker Nodes

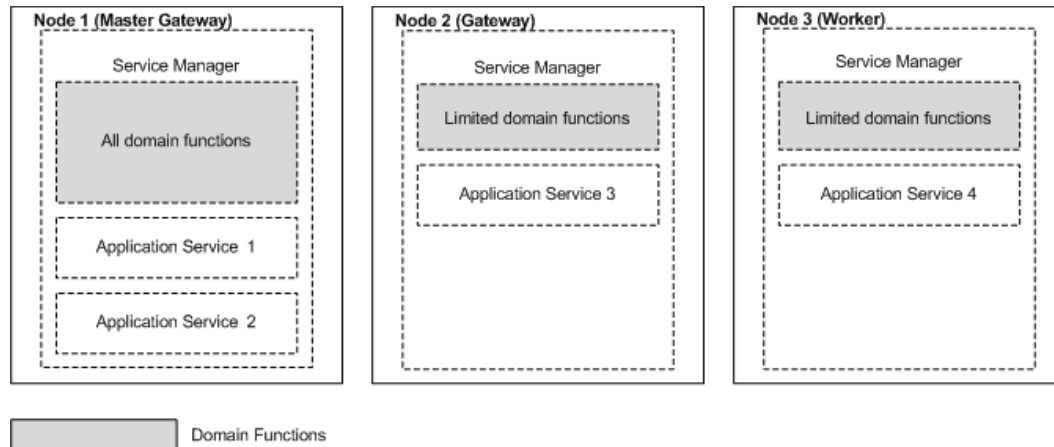
A worker node is any node that you do not configure to serve as a gateway for the domain. A worker node can run application services and perform computations, but it cannot serve as a gateway. The Service Manager performs limited domain functions on a worker node.

Example Domain with Multiple Nodes

This example domain contains three nodes. Each node has both the service and compute roles enabled.

All nodes run the Service Manager. Node 1 is the master gateway node and runs two application services. Node 2 is a back-up gateway node and runs one application service. Node 3 is a worker node and runs one application service. If Node 1 becomes unavailable, Node 2 is elected as the new master gateway node. The Service Manager on Node 2 then performs all domain functions. When Node 1 restarts, it becomes a back-up gateway node and the Service Manager performs limited domain functions.

The following image shows a domain with two gateway nodes and one worker node:



Node Roles

The node role defines the purpose of the node. A node with the service role can run application services. A node with the compute role can perform computations requested by remote application services. A node with both roles can run application services and locally perform computations for those services.

By default, each gateway and worker node has both the service and compute roles enabled. Each node must have at least one role enabled.

You can configure a Data Integration Service grid so that some nodes are dedicated to running application service processes while other nodes are dedicated to performing computations. When you enable only the compute role on a node in a Data Integration Service grid, the node does not have to run the service process. The machine uses all available processing power to run mappings. You can add additional nodes with only the compute role to the grid to increase scalability of Data Integration Service mappings.

For more information about setting up a Data Integration Service grid, see the *Informatica Application Service Guide*.

Service Role

A node with the service role can run application services.

When you enable the service role on a node, the Service Manager supports application services configured to run on that node.

A node requires the service role in the following situations:

- The node is a gateway node.
- The node is configured as a primary or back-up node for an application service.
- The node is assigned to a PowerCenter Integration Service grid or to a Data Integration Service grid and a service process is running on the node.

Compute Role

A node with the compute role can perform computations requested by remote application services.

When a node has the compute role, the Service Manager manages the containers on the node. A container is an allocation of memory and CPU resources. An application service uses the container to remotely perform computations on the node. For example, a Data Integration Service grid includes Node 1 with the service role and Node 2 with the compute role. The Data Integration Service process that runs on Node 1 runs a mapping within a container on Node 2.

A node requires the compute role when the Data Integration Service runs jobs on the node. When the Data Integration Service runs on a single node, the node must have both the service and compute roles. When the Data Integration Service runs on a grid, at least one of the nodes in the grid must have the compute role.

A node does not require the compute role when the Data Integration Service does not run jobs on the node. In this case, you can disable the compute role on the node. However, because the container management function of the Service Manager is a lightweight process, enabling or disabling the compute role does not have a performance impact.

When you disable the compute role on a node, you must specify whether to stop, complete, or abort computations that might be running on the node.

Updating the Node Role

By default, each node has both the service and compute roles. If a node is assigned to a Data Integration Service grid that is configured to run jobs in separate remote processes, you might want to update the node role.

Enable only the service role to dedicate the node to running the Data Integration Service process. Enable only the compute role to dedicate the node to running Data Integration Service mappings.

Note: Before you can disable the service role on a node, you must shut down all application service processes running on the node and remove the node as a primary or back-up node for any application service. You cannot disable the service role on a gateway node.

1. In the Administrator tool, click the **Manage** tab > **Services and Nodes** view.
2. In the Domain Navigator, select a node assigned to a Data Integration Service grid.
3. In the Properties view, click **Edit** for the general properties.
The **Edit General Properties** dialog box appears.
4. Select or clear the service and compute roles to update the node role.
5. Click **OK**.
6. If you disabled the compute role, the **Disable Compute Role** dialog box appears. Perform the following steps:
 - a. Select one of the following modes to disable the compute role:
 - Complete. Allows jobs to run to completion before disabling the role.
 - Stop. Stops all jobs and then disables the role.
 - Abort. Tries to stop all jobs before aborting them and disabling the role.
 - b. Click **OK**.
7. If you updated the role on a node assigned to a Data Integration Service or a Data Integration Service grid, recycle the Data Integration Service.

Viewing Processes on a Node with the Service Role

You can view the status of all application service processes configured to run on a node with the service role. Before you shut down or remove a node, you can view the status of each application service process to determine which service processes you need to disable.

When a node does not have the service role, no application service processes run on the node.

1. In the Administrator tool, click the **Manage** tab > **Services and Nodes** view.
2. In the Domain Navigator, select a node with the service role.
3. In the contents panel, select the **Processes** view.

The view displays the status of each application service process configured to run on the node.

Define and Add Nodes

To create a node, you define the node as a gateway or worker node and then add the node to the domain.

Use either of the following programs to define a node:

Informatica installer

Run the installer on each machine you want to define as a node.

infasetup command line program

Run the `infasetup DefineGatewayNode` or `infasetup DefineWorkerNode` command on each machine that you want to define as a node. You might use `infasetup` to define a node if you decide to move a node from one domain to another domain.

When the Informatica installer or `infasetup` defines a node, the program creates `nodemeta.xml`. This file is the node configuration file for the node. A gateway node uses information in `nodemeta.xml` to connect to the domain configuration database. A worker node uses the information in `nodemeta.xml` to connect to the domain. The file is stored in the following directory on each node:

```
<Informatica installation directory>/isp/config
```

When you define a node using the Informatica installer, the installer adds the node to the domain with both the service and compute roles enabled. When you log in to the Administrator tool, the node appears in the Navigator.

When you define a node with `infasetup`, you must manually add the node to the domain. You can add a node to the domain in the Administrator tool or with the `infacmd isp AddDomainNode` command. When you add the node, you specify the roles to enable on the node.

You can use the Administrator tool to add a node to the domain before you define the node. In this case, the Administrator tool displays a message saying that you need to run the Informatica installer to associate the node with a physical host name and port number. The name that you enter for the node must be the same name that you use when you define the node.

Adding Nodes to the Domain

You can use the Administrator tool to add a node to the domain.

Use the Administrator tool to add a node to the domain in the following situations:

- After you run the `infasetup DefineGatewayNode` or `infasetup DefineWorkerNode` command.

- When you decide to add the node before running the Informatica installer or infasetup command line program to define the node.
1. In the Administrator tool, click the **Manage** tab > **Services and Nodes** view.
 2. In the Domain Navigator, select the folder where you want to add the node. If you do not want the node to appear in a folder, select the domain.
 3. On the Navigator Actions menu, click **New > Node**.
The **Create Node** dialog box appears.
 4. Enter the node name.
The name must be the same node name that you use when you define the node.
 5. If you want to change the folder for the node, click **Browse** and choose a new folder or the domain.
 6. Optionally update the node role.
By default, each node has both the service and compute roles. If a node is assigned to a Data Integration Service grid, you might want to update the node role to dedicate the node to running the Data Integration Service process or to running mappings.
 7. Click **OK**.
If you add a node to the domain before you define the node using the Informatica installer or infasetup, the Administrator tool displays a message saying that you need to run the installer to associate the node with a physical host name and port number.

RELATED TOPICS:

- [“Node Roles” on page 93](#)

Configuring Node Properties

You configure node properties on the Properties view for the node. You can configure properties such as the node roles, error severity level, and minimum and maximum port numbers.

1. In the Administrator tool, click the **Manage** tab > **Services and Nodes** view.
2. In the Domain Navigator, select a node.
3. Click the **Properties** view.
The Properties view displays the node properties in separate sections.
4. In the **Properties** view, click **Edit** for the section that contains the property you want to set.
5. Edit the following properties:

Node Property	Description
Name	Name of the node. The name is not case sensitive and must be unique within the domain. It cannot exceed 128 characters or begin with @. It also cannot contain spaces or the following special characters: ` ~ % ^ * + = { } \ ; : ' " / ? . , < > ! () []
Description	Description of the node. The description cannot exceed 765 characters.

Node Property	Description
Host Name	Host name of the machine represented by the node.
Port	Port number used by the node.
Gateway Node	Indicates whether the node can serve as a gateway. If this property is disabled, then the node is a worker node.
Service Role	Indicates whether the node has the service role. If enabled, application services can run on the node. If disabled, application services cannot run on the node. Disable the property only if the node is assigned to a Data Integration Service grid and you want to dedicate the node to running mappings. Default is enabled.
Compute Role	Indicates whether the node has the compute role. If enabled, the node can perform computations. If disabled, the node cannot perform computations. A node requires the compute role when the Data Integration Service runs jobs on the node. If the Data Integration Service does not run jobs on the node, you can disable the compute role. However, enabling or disabling the compute role does not have a performance impact. Default is enabled.
Backup Directory	Directory to store repository backup files. The directory must be accessible by the node.
Error Severity Level	Level of error logging for the node. These messages are written to the Log Manager application service and Service Manager log files. Set one of the following message levels: <ul style="list-style-type: none"> - ERROR. Writes ERROR code messages to the log. - WARNING. Writes WARNING and ERROR code messages to the log. - INFO. Writes INFO, WARNING, and ERROR code messages to the log. - TRACING. Writes TRACE, INFO, WARNING, and ERROR code messages to the log. - DEBUG. Writes DEBUG, TRACE, INFO, WARNING, and ERROR code messages to the log. Default is WARNING .
Minimum Port Number	Minimum port number used by service processes on the node. To apply changes, restart Informatica services. The default value is the value entered when the node was defined.
Maximum Port Number	Maximum port number used by service processes on the node. To apply changes, restart Informatica services. The default value is the value entered when the node was defined.
CPU Profile Benchmark	Ranks the node's CPU performance against a baseline system. Used by the Load Balancer component of the PowerCenter Integration Service. For example, if the CPU is running 1.5 times as fast as the baseline machine, the value of this property is 1.5. You can calculate the benchmark by clicking Actions > Recalculate CPU Profile Benchmark . The calculation takes approximately five minutes and uses 100% of one CPU on the machine. Or, you can update the value manually. Default is 1.0. Minimum is 0.001. Maximum is 1,000,000. Used in adaptive dispatch mode. Ignored in round-robin and metric-based dispatch modes.

Node Property	Description
Maximum Processes	<p>Maximum number of running session tasks or command tasks allowed for each PowerCenter Integration Service process running on the node. Used by the Load Balancer component of the PowerCenter Integration Service.</p> <p>For example, if you set the value to 5, up to 5 command tasks and 5 session tasks can run at the same time.</p> <p>Set this threshold to a high number, such as 200, to cause the Load Balancer to ignore it. To prevent the Load Balancer from dispatching tasks to this node, set this threshold to 0.</p> <p>Default is 10. Minimum is 0. Maximum is 1,000,000,000.</p> <p>Used in all dispatch modes.</p>
Maximum CPU Run Queue Length	<p>Maximum number of runnable threads waiting for CPU resources on the node. Used by the Load Balancer component of the PowerCenter Integration Service.</p> <p>Set this threshold to a low number to preserve computing resources for other applications. Set this threshold to a high value, such as 200, to cause the Load Balancer to ignore it.</p> <p>Default is 10. Minimum is 0. Maximum is 1,000,000,000.</p> <p>Used in metric-based and adaptive dispatch modes. Ignored in round-robin dispatch mode.</p>
Maximum Memory %	<p>Maximum percentage of virtual memory allocated on the node relative to the total physical memory size. Used by the Load Balancer component of the PowerCenter Integration Service.</p> <p>Set this threshold to a value greater than 100% to allow the allocation of virtual memory to exceed the physical memory size when dispatching tasks. Set this threshold to a high value, such as 1,000, if you want the Load Balancer to ignore it.</p> <p>Default is 150. Minimum is 0. Maximum is 1,000,000,000.</p> <p>Used in metric-based and adaptive dispatch modes. Ignored in round-robin dispatch mode.</p>
Log Collection Directory	<p>The directory that stores the logs for the application service when you run the log aggregator. The directory must be accessible from all the nodes in the domain. If the log collection directory is not accessible by other nodes, the aggregated logs do not appear in the aggregated logs listgrid. The users who run node processes must have read-write permissions on the directory.</p> <p>Configure the log collection directory for the master gateway node in the domain.</p>
Core Dump Directory	<p>The directory that stores the core dump files for the domain processes when you run the log aggregator.</p> <p>Configure the core dump directory for all the nodes in the domain.</p>

- Click **OK**.

Shutting Down and Restarting the Node

Some administrative tasks might require you to shut down a node. For example, you might need to perform maintenance or benchmarking on a machine. You might also need to shut down and restart a node for some configuration changes to take effect. For example, if you change the shared directory for the Log Manager or domain, you must shut down the node and restart it to update the configuration files.

You can shut down a node from the Administrator tool or from the operating system. When you shut down a node, you stop Informatica services and abort all application service processes and computations running on the node.

To restart a node, start Informatica services on the node.

Warning: To avoid loss of data or metadata when you shut down a node, disable all running application service processes in complete mode.

Shutting Down a Node from the Administrator Tool

When you shut down a node from the Administrator tool, you can view all application service processes running on the node.

1. In the Administrator tool, click the **Manage** tab > **Services and Nodes** view.
2. In the Domain Navigator, select a node.
3. In the Navigator **Actions** menu, select **Shutdown Node**.

If the node has the service role, the Administrator tool displays the list of application service processes running on that node.

4. Optionally, choose whether the shutdown is planned or unplanned.
5. Optionally, enter comments about the shutdown.
6. Click **OK** to stop all service processes and shut down the node, or click **Cancel** to cancel the operation.

Starting or Stopping a Node on Windows

On Windows, use the Control Panel to start and stop the Informatica service.

1. Open the Windows Control Panel.
2. Select **Administrative Tools**.
3. Right-click **Services** and select **Run as Administrator**.
4. Right-click the Informatica service.
5. If the service is running, click **Stop**.
If the service is stopped, click **Start**.

Starting or Stopping a Node on UNIX

On UNIX, run `infaservice.sh` to start and stop the Informatica daemon. By default, `infaservice.sh` is installed in the following directory:

```
<InformaticaInstallationDir>/tomcat/bin
```

1. Go to the directory where `infaservice.sh` is located.
2. At the command prompt, enter the following command to start the daemon:

```
infaservice.sh startup
```

Enter the following command to stop the daemon:

```
infaservice.sh shutdown
```

Note: If you use a softlink to specify the location of `infaservice.sh`, set the `INFA_HOME` environment variable to the location of the Informatica installation directory.

Removing the Node Association

You can remove the host name and port number associated with a node. When you remove the node association, the node remains in the domain, but it is not associated with a host machine.

To associate a different host machine with the node, you must run the installation program or `infasetup DefineGatewayNode` or `infasetup DefineWorkerNode` command on the new host machine, and then restart the node on the new host machine.

1. In the Administrator tool, click the **Manage** tab > **Services and Nodes** view.
2. In the Navigator, select a node.
3. In the **Services and Nodes** view **Actions** menu, select **Remove Node Association**.

Removing a Node

When you remove a node from a domain, it is no longer visible in the Navigator. If the node is running when you remove it, the node shuts down and aborts all application service processes.

Note: To avoid loss of data or metadata when you remove a node, disable all running application service processes in complete mode.

1. In the Administrator tool, click the **Manage** tab > **Services and Nodes** view.
2. In the Domain Navigator, select a node.
3. In the Navigator **Actions** menu, select **Delete**.
4. In the warning message that appears, click **OK**.

CHAPTER 7

High Availability

This chapter includes the following topics:

- [High Availability Overview, 101](#)
- [Resilience, 102](#)
- [Restart and Failover, 105](#)
- [Recovery, 106](#)
- [Configuration for a Highly Available Domain, 107](#)
- [Troubleshooting High Availability, 111](#)

High Availability Overview

High availability refers to the uninterrupted availability of computer system resources. In an Informatica domain, high availability eliminates a single point of failure and provides minimal service interruption in the event of failure. When you configure high availability for a domain, the domain can continue running despite temporary network, hardware, or service failures.

The following high availability components make services highly available in an Informatica domain:

- **Resilience.** An Informatica domain can tolerate temporary connection failures until either the resilience timeout expires or the failure is fixed.
- **Restart and failover.** A process can restart on the same node or on a backup node after the process becomes unavailable.
- **Recovery.** Operations can complete after a service is interrupted. After a service process restarts or fails over, it restores the service state and recovers operations.

When you plan a highly available Informatica environment, configure high availability for both the internal Informatica components and systems that are external to Informatica. Internal components include the domain, application services, application clients, and command line programs. External systems include the network, hardware, database management systems, FTP servers, message queues, and shared storage.

High availability features for the Informatica environment are available based on your license.

Example

As you open a mapping in the PowerCenter Designer workspace, the PowerCenter Repository Service becomes unavailable and the request fails. The domain contains multiple nodes for failover and the PowerCenter Designer is resilient to temporary failures.

The PowerCenter Designer tries to establish a connection to the PowerCenter Repository Service within the resilience timeout period. The PowerCenter Repository Service fails over to another node because it cannot restart on the same node.

The PowerCenter Repository Service restarts within the resilience timeout period, and the PowerCenter Designer reestablishes the connection.

After the PowerCenter Designer reestablishes the connection, the PowerCenter Repository Service recovers from the failed operation and fetches the mapping into the PowerCenter Designer workspace.

Resilience

The domain tolerates temporary connection failures between application clients, application services, and nodes.

A temporary connection failure might occur because an application service process fails or because of a network failure. When a temporary connection failure occurs, the Service Manager tries to reestablish connections between the application clients, application services, and nodes.

Application Client Resilience

The application clients try to reconnect to application services when a temporary connection failure occurs.

Based on your license, the following application clients are resilient to the services that they connect to:

Developer Tool Client

The Developer tool client tries to reconnect to the Data Integration Service or the Data Integration Service grid when a temporary network failure occurs.

If a job is running and the Developer tool cannot reconnect to the Data Integration Service or the Data Integration Service grid within the re-connection timeout period, the Developer tool does not resubmit the job to a Data Integration Service or a Data Integration Service grid on a different node. The Developer tool client fails the job.

PowerCenter Client

The PowerCenter Client tries to reconnect to the PowerCenter Repository Service and the PowerCenter Integration Service when a temporary network failure occurs.

If you perform a PowerCenter Client action that requires connection to the repository while the PowerCenter Client is trying to reestablish the connection, the PowerCenter Client prompts you to try the operation again after the PowerCenter Client reestablishes the connection. If the PowerCenter Client is unable to reestablish the connection during the resilience timeout period, the PowerCenter Client prompts you to reconnect to the repository manually.

Command line programs

Command line programs try to reconnect to the domain or an application service when a temporary network failure occurs while a command line program is running.

If the command line program is running on a Data Integration Service or a Data Integration Service grid and the command line program cannot reconnect to the Data Integration Service or the Data Integration Service grid within the re-connection timeout period, the command line program does not resubmit the job to a Data Integration Service or a Data Integration Service grid on a different node. The command line program fails the command.

Example PowerCenter Client Resilience to Application Services

There is a network connection loss of 120 seconds between the PowerCenter Workflow Monitor and the PowerCenter Repository Service when a developer is monitoring a workflow. The PowerCenter client, Workflow Monitor has a 60 second resilience timeout and the PowerCenter Repository Service has a resilience timeout of 180 seconds.

The Developer does not notice the loss of connection and he is unaffected by the 120 seconds connection loss. However, the following messages appear in the **Notifications** tab on the PowerCenter Workflow Monitor:

```
Repository Service notifications are enabled.  
DATE TIME-[REP_55101] Connection to the Repository Service [Repository_Service_Name] is  
broken.  
DATE TIME-[REP_55114] Reconnecting to the Repository Service [Repository_Service_Name].  
The resilience time is 180 seconds.  
DATE TIME-Reconnected to Repository Service [Repository_Service_Name] successfully.
```

Application Service Resilience

Some application services try to reconnect to application services, application clients, and external components when a temporary connection failure occurs.

Based on your license, the following application services are resilient to the temporary connection failure of their clients:

Data Integration Service

The Data Integration Service is resilient to temporary connection failures to other services and the Developer tool client.

PowerCenter Integration Service

The PowerCenter Integration Service is resilient to temporary connection failures to other services, the PowerCenter client, and external components such as databases and FTP servers.

PowerCenter Repository Service

The PowerCenter Repository Service is resilient to temporary connection failures to other services, such as the PowerCenter Integration Service. It is also resilient to the temporary connection failures to the repository database.

Node Resilience

When a domain contains multiple nodes, the nodes are resilient to temporary failures in communication from other nodes in the domain.

Nodes are resilient to the following temporary connection failures:

A non-master gateway node becomes unavailable.

Every node in the domain sends a communication signal to the master gateway node at periodic intervals of 15 seconds. For nodes with the service role, the communication includes a list of application services running on the node.

All nodes have a resilience timeout of 90 seconds. If a node fails to connect to the master gateway node within the resilience timeout period, the master gateway node marks the node unavailable. If the node that fails to connect has the service role, the master gateway node also reassigns its application services to a back-up node. This ensures that services on a node continue to run despite node failures.

The master gateway node becomes unavailable.

You can configure more than one node to serve as a gateway. If the master gateway node becomes unavailable, the Service Managers on the other gateway nodes elect another master gateway node.

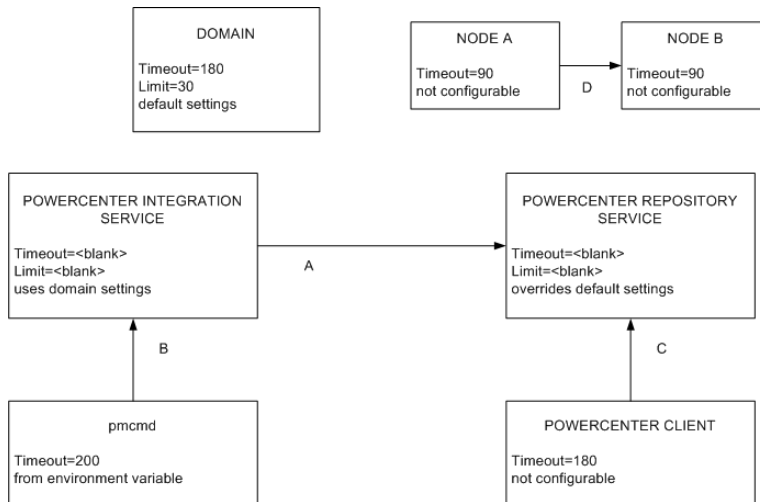
If you configure one node to serve as the gateway and the node becomes unavailable, all other nodes shut down.

Example Resilience Timeout Configuration

Some resilience timeout values are default and others can be configured or overwritten.

You can use the resilience timeout and limit on resilience timeout configured for the domain for PowerCenter application services if you do not set it for the application service. Command line programs use the service resilience timeout. If the service limit on resilience timeout is smaller than the resilience timeout for the connecting client, the client uses the services limit as the resilience timeout.

The following figure shows some sample connections and resilience configurations in a domain with PowerCenter application services:



The following table describes the resilience timeout and the limits shown in the figure above:

	Connect From	Connect To	Description
A	PowerCenter Integration Service	PowerCenter Repository Service	The PowerCenter Integration Service can spend up to 30 seconds to connect to the PowerCenter Repository Service, based on the domain resilience timeout. It is not bound by the PowerCenter Repository Service limit on resilience timeout of 60 seconds.
B	<i>pmcmd</i>	PowerCenter Integration Service	<i>pmcmd</i> is bound by the PowerCenter Integration Service limit on resilience timeout of 180 seconds, and it cannot use the 200 second resilience timeout configured in INFA_CLIENT_RESILIENCE_TIMEOUT.

	Connect From	Connect To	Description
C	PowerCenter Client	PowerCenter Repository Service	The PowerCenter Client is bound by the PowerCenter Repository Service limit on resilience timeout of 60 seconds. It cannot use the default resilience timeout of 180 seconds.
D	Node A	Node B	Node A can spend up to 90 seconds to connect to Node B. The Service Managers on Node A and Node B use the default node resilience timeout of 90 seconds.

Restart and Failover

To maximize operation time in the event of a failure, the Informatica domain can restart or fail over processes to another node.

The Service Manager on the master gateway node accepts application service request and manages the domain. If a master gateway node is not available, the domain shuts down. Configure the domain to failover to another node by configuring multiple gateway nodes.

Based on your license, you can also configure backup nodes for application services. The Service Manager can restart or failover the following application services if a failure occurs:

- Data Integration Service
- Model Repository Service
- PowerCenter Integration Service
- PowerCenter Repository Service
- PowerExchange Listener Service
- PowerExchange Logger Service
- Resource Manager Service

Domain Failover

The Service Manager on the master gateway node accepts service requests and manages the domain and services in the domain. The domain can failover to another node when the domain has multiple gateway nodes. Configure multiple gateway nodes to prevent domain shutdown when the master gateway node is unavailable.

The master gateway node maintains a connection to the domain configuration repository. If the domain configuration repository becomes unavailable, the master gateway node tries to reconnect when a user performs an operation. If the master gateway node cannot connect to the domain configuration repository, the master gateway node may shut down.

If the domain has multiple gateway nodes and the master gateway node becomes unavailable, the Service Managers on the other gateway nodes elect another master gateway node. The domain tries to connect to the domain configuration repository with each gateway node. If none of the gateway nodes can connect, the domain shuts down and all domain operations fail. When a master gateway fails over, the client tools retrieve information about the alternate domain gateways from the domains.infa file.

Note: Application services running on the master gateway node will not fail over when another master gateway node is elected unless the application service has a backup node configured.

Application Service Restart and Failover

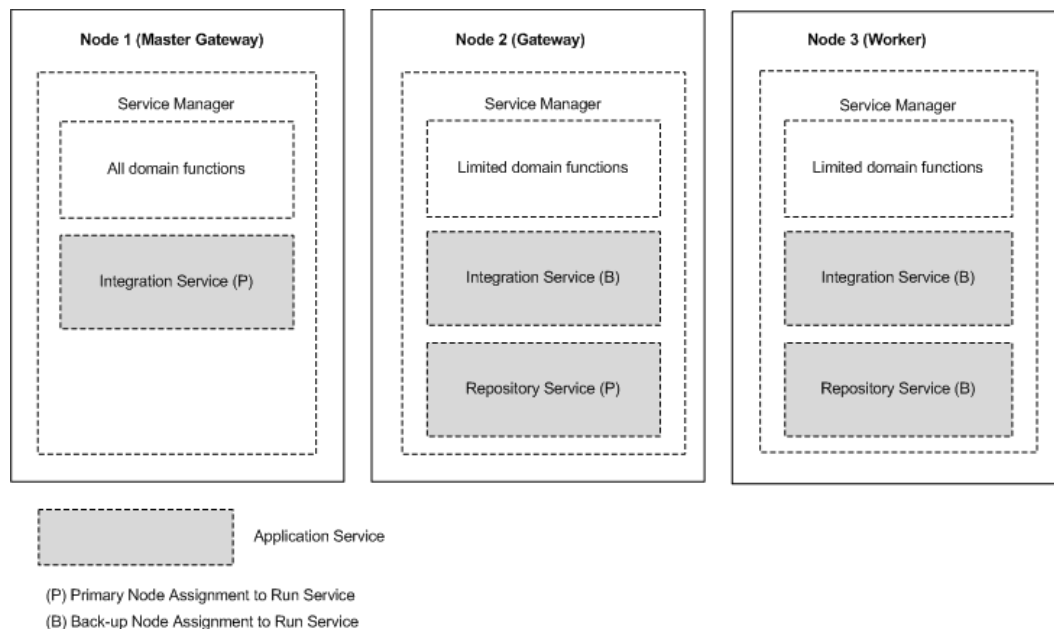
If an application service process becomes unavailable, the Service Manager can restart the application service or fail it over to a back-up node. When the Service Manager fails over an application service, it starts the service on another node that the service is configured to run on.

The following situations describe how the Service Manager restarts or fails over an application service:

- If the primary node running the service process becomes unavailable, the service fails over to a back-up node. The primary node might be unavailable if it shuts down or if the connection to the node becomes unavailable.
- If the primary node running the service process is available, the domain tries to restart the process based on the restart options configured in the domain properties. If the process does not restart, the Service Manager may mark the process as failed. The service then fails over to a back-up node and starts another process. If the Service Manager marks the process as failed, the administrator must enable the process after addressing any configuration problem.

If a service process fails over to a back-up node, it does not fail back to the primary node when the node becomes available. You can disable the service process on the back-up node to cause it to fail back to the primary node.

The following image shows how you can configure primary and back-up nodes for an application service:



Recovery

Recovery is the completion of operations after an interrupted service is restored. The state of operation for a service contains information about the service process.

Based on your license, the following components can recover after an interrupted service is restored:

Service Manager

The Service Manager for each node in the domain maintains the state of service processes running on that node. If the master gateway shuts down, the newly elected master gateway collects the state information from each node to restore the state of the domain.

PowerCenter Repository Service

The PowerCenter Repository Service maintains the state of operation in the PowerCenter repository. The state of operation includes information about repository locks, requests in progress, and connected clients. After restart or failover, the PowerCenter Repository Service can recover operations from the point of interruption.

PowerCenter Integration Service

The PowerCenter Integration Service maintains the state of operation in the shared storage configured for the service. The state of operation includes information about scheduled, running, and completed tasks for the service.

The PowerCenter Integration Service maintains PowerCenter session and workflow state of operation based on the recovery strategy you configure for the session and workflow. When the PowerCenter Integration Service restarts or fails over a service process, it can automatically recover interrupted workflows that are configured for recovery.

Data Integration Service

The Data Integration Service maintains the state of operation in the Model repository. The state of operation includes the state of the workflow and workflow tasks and the values of the workflow variables and parameters during the interrupted workflow instance.

When a Data Integration Service process restarts or fails over a service process, you can manually restart interrupted workflows that are enabled for workflow recovery.

Configuration for a Highly Available Domain

To minimize system downtime, configure Informatica domain components to be highly available.

You can configure the following Informatica domain components to be highly available:

Domain

One node in the domain acts as a gateway to receive service requests from clients and routes them to the appropriate service and node. To prevent domain shutdown when the master gateway node is unavailable, configure more than one gateway node.

Nodes

Informatica services are processes that run on each node. You can configure Informatica services to restart automatically if it terminates unexpectedly.

Application Services

The application services run on nodes in the Informatica domain.

Based on your license, you can configure the following high availability features for application services:

- To minimize the application service downtime, configure backup nodes for application services.
- To specify the resilience period for application services, review default settings and configure resilience timeout periods for application services.

- To ensure PowerCenter Integration Service failover and recovery, configure the PowerCenter Integration Service to store process state information on a POSIX compliant shared file system or in a database.

Application Clients

Application clients provide access to Informatica functionality, and they run on user machines. Application clients send requests to the Service Manager or application services.

You can configure resilience timeout periods for command line programs. You cannot configure a PowerCenter client resilience timeout.

External Systems

Use highly available versions of external systems such as source and target databases, message queues, and FTP servers.

Network

Make the network highly available by configuring redundant components such as routers, cables, and network adapters.

Application Service Resilience Configuration

When a temporary network failure occurs, application services try to reconnect to other application services for the duration of the resilience timeout. You can configure the resilience timeout for application services.

When an application service connects to another application service in the domain, the service that initiates the connection is a client of the other service.

You can configure application service resilience timeouts for the following application services:

PowerCenter Application Services

You can configure the resilience timeout and resilience timeout limits in the advanced properties of the PowerCenter Integration Service and PowerCenter Repository Service. The resilience timeout for application services that connects to a PowerCenter Integration Service or PowerCenter Repository Service is determined by one of the following values:

- The service **Resilience Timeout** property. You can configure the resilience timeout for the service in the service properties. To disable resilience for a service, set the resilience timeout to 0.
- The domain **Resilience Timeout** property. To use the resilience timeout configured for the domain, set the resilience timeout for the service to blank.
- The service **Limit on Resilience Timeout** property. If the service limit on resilience timeout is smaller than the resilience timeout for the connecting client, the client uses the limit as the resilience timeout. To use the limit on resilience timeout configured for the domain, set the service resilience limit to blank.
- The domain **Limit on Resilience Timeout** property. To use the resilience timeout configured for the domain, set the limit on resilience timeout for the service to blank.

You can configure the resilience timeout for the SAP BW Service in the general properties for the service. The SAP BW Service resilience timeout property is called the **Retry Period**.

Note: A client cannot be resilient to service interruptions if you disable the service in the Administrator tool. If you disable the service process, the client is resilient to the interruption in service.

Application Service Failover Configuration

Based on your license, you can configure backup nodes so that application services can failover to another node when the primary node fails. Configure backup nodes when you create or update an application service.

When you configure a backup node, verify that the node has access to run-time files that each application service requires to process data integration tasks such as workflows and mappings. For example, a workflow might require parameter files, input files, or output files.

PowerCenter Integration Service Failover and Recovery Configuration

During failover and recovery, the PowerCenter Integration Service needs to access state of operation files and process state information.

The state of operation files store the state of each workflow and session operation. The PowerCenter Integration Service always stores the state of each workflow and session operation in files in the \$PMStorageDir directory of the PowerCenter Integration Service process.

Process state information includes information about which node was running the master PowerCenter Integration Service process and which node was running each session. You can configure the PowerCenter Integration Service to store process state information on a cluster file system or in the PowerCenter repository database.

Store High Availability Persistence on a Cluster File System

By default, the PowerCenter Integration Service stores process state information along with the state of operation files in the \$PMStorageDir directory of the Integration Service process. You must configure the \$PMStorageDir directory for each PowerCenter Integration Service process to use the same directory on a cluster file system.

Nodes that run the PowerCenter Integration Service must be on the same cluster file system so that they can share resources. Also, nodes within a cluster must be on the cluster file system's heartbeat network. Use a highly available cluster file system that is configured for I/O fencing. The hardware requirements and configuration of an I/O fencing solution are different for each file system.

The following cluster file systems are certified by Informatica for use for PowerCenter Integration Service failover and session recovery:

Storage Array Network

- Veritas Cluster Files System (VxFS)

- IBM General Parallel File System (GPFS)

Network Attached Storage using NFS v3 protocol

- EMC UxFS hosted on an EMV Celerra NAS appliance

- NetApp WAFL hosted on a NetApp NAS appliance

Contact the file system vendors directly to evaluate which file system matches your requirements.

Store High Availability Persistence in a Database

You can configure the PowerCenter Integration Service to store process state information in database tables. When you configure the PowerCenter Integration Service to store process state information in a database, the service still stores the state of each workflow and session operation in files in the \$PMStorageDir directory. You can configure the \$PMStorageDir directory to use a POSIX compliant shared file system. You do not need to use a cluster file system.

Configure the PowerCenter Integration Service to store process state information in database tables in the advanced properties. The PowerCenter Integration Service stores process state information in persistent database tables in the associated PowerCenter repository database.

During failover, automatic recovery of workflows resume when the service process can access the database tables.

Command Line Program Resilience Configuration

You can configure the resilience timeout that command line programs use to perform domain and service operations.

When you use the `infacmd`, `pmcmd`, or `pmrep` command line programs to connect to the domain or an application service the resilience timeout is determined by the command line option, an environment variable, or the default resilience timeout.

Use the following guidelines when you configure command line program resilience:

Command line option

You can set the resilience timeout for `infacmd` by using the `-ResilienceTimeout` command line option each time you run a command. You can set the resilience timeout for `pmcmd` by using the `-timeout` command line option each time you run a command. When you use `pmrep` connect to connect to a repository, you can use the `-t` command line option to set the resilience timeout for `pmrep` commands that use the connection.

Environment variable.

If you do not set the timeout option in the `infacmd` and `pmcmd` command line syntax, the `infacmd` and `pmcmd` command line programs use the value of the environment variable `INFA_CLIENT_RESILIENCE_TIMEOUT` that is configured on the client machine. If you do not set the timeout option when you use `pmrep` connect to connect to the repository, `pmrep` commands use the value of the environment variable `INFA_CLIENT_RESILIENCE_TIMEOUT` that is configured on the client machine.

Default value

If you do not use the command line option or the environment variable, the `pmcmd` and `pmrep` command line program uses the default resilience timeout of 180 seconds. If you do not use the command line option or the environment variable, the `infacmd` command line program uses the value of the domain **Service Level Timeout** property as the default resilience timeout.

Limit on timeout

If the limit on resilience timeout for the PowerCenter Integration Service or the PowerCenter Repository Service is smaller than the command line resilience timeout, the command line program uses the limit as the resilience timeout.

Note: PowerCenter does not provide resilience for a repository client when the PowerCenter Repository Service is running in exclusive mode.

Domain Failover Configuration

You can define multiple gateway nodes to prevent domain shutdown when the master gateway node is unavailable.

The first time that you install Informatica services, you create one gateway node. After you install Informatica, you can define additional gateway nodes. To define a gateway node, add a gateway node to the domain or configure a worker node to serve as a gateway node.

Node Restart Configuration

The Informatica services run the Service Manager on all nodes in the domain. You can configure the Informatica services to start automatically when a node terminates unexpectedly and restarts.

To restart the Informatica services when a node restarts, complete the following steps:

- In a UNIX environment, you can create a script to automatically start the Informatica services when the node starts.
- In a Windows environment, go to the Control Panel and configure the Informatica services to start automatically.

You can configure restart for all nodes, regardless of node type or node role.

Troubleshooting High Availability

The solutions to the following situations might help you with high availability.

[I am not sure where to look for status information regarding client connections to the PowerCenter repository.](#)

In PowerCenter Client applications such as the PowerCenter Designer and the PowerCenter Workflow Manager, an error message appears if the connection cannot be established during the timeout period. Detailed information about the connection failure appears in the Output window. If you are using *pmrep*, the connection error information appears at the command line. If the PowerCenter Integration Service cannot establish a connection to the repository, the error appears in the PowerCenter Integration Service log, the workflow log, and the session log.

[I entered the wrong connection string for an Oracle database. Now I cannot enable the PowerCenter Repository Service even though I edited the PowerCenter Repository Service properties to use the right connection string.](#)

You need to wait for the database resilience timeout to expire before you can enable the PowerCenter Repository Service with the updated connection string.

[I have the high availability option, but my FTP server is not resilient when the network connection fails.](#)

The FTP server is an external system. To achieve high availability for FTP transmissions, you must use a highly available FTP server. For example, Microsoft IIS 6.0 does not natively support the restart of file uploads or file downloads. File restarts must be managed by the client connecting to the IIS server. If the transfer of a file to or from the IIS 6.0 server is interrupted and then reestablished within the client resilience timeout period, the transfer does not necessarily continue as expected. If the write process is more than half complete, the target file may be rejected.

[I have the high availability option, but the Informatica domain is not resilient when machines are connected through a network switch.](#)

If you are using a network switch to connect machines in the domain, use the auto-select option for the switch.

CHAPTER 8

Connections

This chapter includes the following topics:

- [Connections Overview, 112](#)
- [Connection Management, 112](#)
- [Pass-through Security, 115](#)
- [Pooling Properties in Connection Objects, 117](#)

Connections Overview

A connection is a repository object that defines a connection in the domain configuration repository.

The Data Integration Service uses database connections to process jobs for the Developer tool and the Analyst tool. Jobs include mappings, data profiles, scorecards, and SQL data services.

You can create and manage connections in the Administrator tool, the Developer tool, and the Analyst tool.

The tasks that you can perform in each tool depend on the tool that you use. For example, you can create an SAP NetWeaver connection in the Developer tool and manage it in the Administrator tool, but you cannot create or manage it in the Analyst tool.

Note: These connections are independent of the connections that you create in the PowerCenter Workflow Manager.

Connection Management

After you create a connection, you can view the connection, configure connection properties, and delete the connection.

After you create a connection, you can perform the following actions on the connection:

Configure connection pooling.

Configure connection pooling to optimize processing for the Data Integration Service. Connection pooling is a framework to cache connections.

View connection properties.

View the connection properties through the **Connections** view on the **Manage** tab.

Edit the connection.

You can change the connection name and the description. You can also edit connection details such as the user name, password, and connection strings. When you update a database connection that has connection pooling disabled, all updates take effect immediately.

The Data Integration Service identifies connections by the connection ID instead of the connection name. When you rename a connection, the Developer tool and the Analyst tool update the jobs that use the connection.

Deployed applications and parameter files identify a connection by name, not by connection ID. Therefore, when you rename a connection, you must redeploy all applications that use the connection. You must also update all parameter files that use the connection parameter.

Delete the connection.

When you delete a connection, objects that use the connection are no longer valid. If you accidentally delete a connection, you can re-create it by creating another connection with the same connection ID as the deleted connection.

Refresh the connections list.

You can refresh the connections list to see the latest list of connections for the domain. Refresh the connections list after a user adds, deletes, or renames a connection in the Developer tool or the Analyst tool.

Creating a Connection

In the Administrator tool, you can create relational database, social media, and file systems connections.

1. In the Administrator tool, click the **Manage** tab.
2. Click the **Connections** view.
3. In the Navigator, select the domain.
4. In the Navigator, click **Actions > New > Connection**.
The **New Connection** dialog box appears.
5. In the **New Connection** dialog box, select the connection type, and then click **OK**.
The **New Connection** wizard appears.
6. Enter the connection properties.
The connection properties that you enter depend on the connection type. Click **Next** to go to the next page of the **New Connection** wizard.
7. When you finish entering connection properties, you can click **Test Connection** to test the connection.
8. Click **Finish**.

Refreshing the Connections List

Refresh the connections list to see the latest list of connections in the domain.

The Administrator tool displays the latest list of connections when you start the Administrator tool. You might want to refresh the connections list when a user adds, deletes, or renames a connection in the Developer tool or the Analyst tool.

1. In the Administrator tool, click the **Manage** tab.
2. Click the **Connections** view.

The Navigator shows all connections in the domain.

3. In the Navigator, select the domain.
4. Click **Actions > Refresh**.

Viewing a Connection

View connections in the Administrator tool.

1. In the Administrator tool, click the **Manage** tab.
2. Click the **Connections** view.

The Navigator shows all connections in the domain.

3. In the Navigator, select the domain.

The contents panel shows all connections for the domain.

4. To filter the connections that appear in the contents panel, enter filter criteria and click the Filter button.

The contents panel shows the connections that meet the filter criteria.

5. To remove the filter criteria, click the Reset Filters button.

The contents panel shows all connections in the domain.

6. To sort the connections, click in the header for the column by which you want to sort the connections.

By default, connections are sorted by name.

7. To add or remove columns from the contents panel, right-click a column header.

If you have Read permission on the connection, you can view the data in the **Created By** column. Otherwise, this column is empty.

8. To view the connection details, select a connection in the Navigator.

The contents panel shows the connection details.

Configuring Pooling for a Connection

Configure pooling for a connection in the Administrator tool.

1. In the Administrator tool, click the **Manage** tab.
2. Click the **Connections** view.
3. In the Domain Navigator, select a connection.

The contents panel shows the connection properties.

4. In the contents panel, click the **Pooling** view.
5. In the **Pooling Properties** area, click **Edit**.

The **Edit Pooling Properties** dialog box appears.

6. Edit the pooling properties and click **OK**.

Editing and Testing a Connection

In the Administrator tool, you can edit connections that you created in the Administrator tool, the Analyst tool, the Developer tool, or by running the `infacmd isp CreateConnection` command. You can test relational database connections.

1. In the Administrator tool, click the **Manage** tab.

2. Click the **Connections** view.
The Navigator shows all connections in the domain.
3. In the Navigator, select a connection.
The contents panel shows properties for the connection.
4. In the contents panel, select the **Properties** view or the **Pooling** view.
5. To edit properties in a section, click **Edit**.
Edit the properties and click **OK**.
Note: If you change a connection name, you must redeploy all applications that use the connection. You must also update all parameter files that use the connection parameter.
6. To test a database connection, select the connection in the Navigator.
Click **Actions > Test Connection** on the **Manage** tab.
Note: The **Test Connection** button tests the connection string of the metadata access properties and not the data access properties.
A message box displays the result of the test.

Deleting a Connection

You can delete a database connection in the Administrator tool.

When you delete a connection in the Administrator tool, you also delete it from the Developer tool and the Analyst tool.

1. In the Administrator tool, click the **Manage** tab.
2. Click the **Connections** view.
The Navigator shows all connections in the domain.
3. In the Navigator, select a connection.
4. In the Navigator, click **Actions > Delete**.

Pass-through Security

Pass-through security is the capability to connect to an SQL data service or an external source with the client user credentials instead of the credentials from a connection object.

Users might have access to different sets of data based on the job in the organization. Client systems restrict access to databases by the user name and the password. When you create an SQL data service, you might combine data from different systems to create one view of the data. However, when you define the connection to the SQL data service, the connection has one user name and password.

If you configure pass-through security, you can restrict users from some of the data in an SQL data service based on their user name. When a user connects to the SQL data service, the Data Integration Service ignores the user name and the password in the connection object. The user connects with the client user name or the LDAP user name.

A web service operation mapping might need to use a connection object to access data. If you configure pass-through security and the web service uses WS-Security, the web service operation mapping connects to a source using the user name and password provided in the web service SOAP request.

Configure pass-through security for a connection in the connection properties of the Administrator tool or with `infacmd dis UpdateServiceOptions`. You can set pass-through security for connections to deployed applications. You cannot set pass-through security in the Developer tool. Only SQL data services and web services recognize the pass-through security configuration.

Example

An organization combines employee data from multiple databases to present a single view of employee data in an SQL data service. The SQL data service contains data from the Employee and Compensation databases. The Employee database contains name, address, and department information. The Compensation database contains salary and stock option information.

A user might have access to the Employee database but not the Compensation database. When the user runs a query against the SQL data service, the Data Integration Service replaces the credentials in each database connection with the user name and the user password. The query fails if the user includes salary information from the Compensation database.

Pass-Through Security with Data Object Caching

To use data object caching with pass-through security, you must enable caching in the pass-through security properties for the Data Integration Service.

When you deploy an SQL data service or a web service, you can choose to cache the logical data objects in a database. You must specify the database in which to store the data object cache. The Data Integration Service validates the user credentials for access to the cache database. If a user can connect to the cache database, the user has access to all tables in the cache. The Data Integration Service does not validate user credentials against the source databases when caching is enabled.

For example, you configure caching for the EmployeeSQLDS SQL data service and enable pass-through security for connections. The Data Integration Service caches tables from the Compensation and the Employee databases. A user might not have access to the Compensation database. However, if the user has access to the cache database, the user can select compensation data in an SQL query.

When you configure pass-through security, the default is to disallow data object caching for data objects that depend on pass-through connections. When you enable data object caching with pass-through security, verify that you do not allow unauthorized users access to some of the data in the cache. When you enable caching for pass-through security connections, you enable data object caching for all pass-through security connections.

Adding Pass-Through Security

Enable pass-through security for a connection in the connection properties. Enable data object caching for pass-through security connections in the pass-through security properties of the Data Integration Service.

1. Select a connection.
2. Click the **Properties** view.
3. Edit the connection properties.
The **Edit Connection Properties** dialog box appears.
4. To choose pass-through security for the connection, select the **Pass-through Security Enabled** option.
5. Optionally, select the Data Integration Service for which you want to enable object caching for pass-through security.
6. Click the **Properties** view.
7. Edit the pass-through security options.

The **Edit Pass-through Security Properties** dialog box appears.

8. Select **Allow Caching** to allow data object caching for the SQL data service or web service. This applies to all connections.
9. Click **OK**.

You must recycle the Data Integration Service to enable caching for the connections.

Pooling Properties in Connection Objects

You can edit connection pooling properties in the **Pooling** view for a database connection.

The number of connection pool libraries depends on the number of running Data Integration Service processes or DTM processes. Each Data Integration Service process or DTM process maintains its own connection pool library. The values of the pooling properties are for each connection pool library.

For example, if you set maximum connections to 15, then each connection pool library can have a maximum of 15 idle connections in the pool. If the Data Integration Service runs jobs in separate local processes and three DTM processes are running, then you can have a maximum of 45 idle connection instances.

To decrease the total number of idle connection instances, set the minimum number of connections to 0 and decrease the maximum idle time for each database connection.

The following list describes database connection pooling properties that you can edit in the **Pooling** view for a database connection:

Enable Connection Pooling

Enables connection pooling. When you enable connection pooling, each connection pool retains idle connection instances in memory. To delete the pools of idle connections, you must restart the Data Integration Service.

If connection pooling is disabled, the DTM process or the Data Integration Service process stops all pooling activity. The DTM process or the Data Integration Service process creates a connection instance each time it processes a job. It drops the instance when it finishes processing the job.

Default is enabled for DB2 for i5/OS, DB2 for z/OS, IBM DB2, Microsoft SQL Server, Oracle, and ODBC connections. Default is disabled for Adabas, IMS, Sequential, and VSAM connections.

Minimum # of Connections

The minimum number of idle connection instances that a pool maintains for a database connection after the maximum idle time is met. Set this value to be equal to or less than the maximum number of idle connection instances. Default is 0.

Maximum # of Connections

The maximum number of idle connection instances that a pool maintains for a database connection before the maximum idle time is met. Set this value to be more than the minimum number of idle connection instances. Default is 15.

Maximum Idle Time

The number of seconds that a connection instance that exceeds the minimum number of connection instances can remain idle before the connection pool drops it. The connection pool ignores the idle time when the connection instance does not exceed the minimum number of idle connection instances. Default is 120.

CHAPTER 9

Connection Properties

This chapter includes the following topics:

- [Connection Properties Overview, 119](#)
- [Adabas Connection Properties, 120](#)
- [Amazon Redshift Connection Properties, 122](#)
- [Amazon S3 Connection Properties, 124](#)
- [Blockchain Connection Properties, 126](#)
- [Cassandra Connection Properties, 128](#)
- [Confluent Kafka Connection, 129](#)
- [Databricks Connection Properties, 131](#)
- [Greenplum Connection Properties, 133](#)
- [Google Analytics Connection Properties, 135](#)
- [Google BigQuery Connection Properties, 135](#)
- [Google Cloud Spanner Connection Properties, 137](#)
- [Google Cloud Storage Connection Properties, 138](#)
- [Google PubSub Connection Properties, 139](#)
- [Hadoop Connection Properties, 140](#)
- [HBase Connection Properties, 146](#)
- [HDFS Connection Properties, 147](#)
- [HBase Connection Properties for MapR-DB, 149](#)
- [Hive Connection Properties, 149](#)
- [HTTP Connection Properties, 153](#)
- [IBM DB2 Connection Properties, 155](#)
- [IBM DB2 for i5/OS Connection Properties, 158](#)
- [IBM DB2 for z/OS Connection Properties, 161](#)
- [IMS Connection Properties, 164](#)
- [JDBC Connection Properties, 166](#)
- [JDBC V2 Connection Properties, 169](#)
- [JD Edwards EnterpriseOne Connection Properties, 171](#)
- [Kafka Connection Properties, 172](#)
- [Kudu Connection Properties, 175](#)
- [LDAP Connection Properties, 176](#)

- [Microsoft Azure Blob Storage Connection Properties, 177](#)
- [Microsoft Azure Cosmos DB SQL API Connection Properties, 178](#)
- [Microsoft Azure Data Lake Storage Gen1 Connection Properties, 179](#)
- [Microsoft Azure Data Lake Storage Gen2 Connection Properties, 180](#)
- [Microsoft Azure SQL Data Warehouse Connection Properties, 181](#)
- [MS SQL Server Connection Properties, 182](#)
- [Netezza Connection Properties, 186](#)
- [OData Connection Properties, 187](#)
- [ODBC Connection Properties, 188](#)
- [Oracle Connection Properties, 189](#)
- [Salesforce Connection Properties, 192](#)
- [Salesforce Marketing Cloud Connection Properties, 193](#)
- [SAP Connection Properties, 194](#)
- [Sequential Connection Properties, 197](#)
- [Snowflake Connection Properties, 199](#)
- [Teradata Parallel Transporter Connection Properties, 200](#)
- [Tableau Connection Properties, 202](#)
- [Tableau V3 Connection Properties, 203](#)
- [Twitter Streaming Connection Properties, 204](#)
- [VSAM Connection Properties, 205](#)
- [Web Services Connection Properties, 208](#)
- [Identifier Properties in Database Connections, 209](#)

Connection Properties Overview

Connection properties enable the Informatica client to connect to data sources.

This chapter contains connection properties for each of the connections you can create and manage through Informatica clients.

Adabas Connection Properties

Use an Adabas connection to access an Adabas database. The Adabas connection is a mainframe database type connection. You create an Adabas connection in the Developer tool. You can manage an Adabas connection in the Administrator tool or the Developer tool.

The following table describes Adabas connection properties:

Option	Description
Location	Node name for the location of the PowerExchange Listener that connects to Adabas. The node name is defined in the first parameter of the NODE statement in the PowerExchange dbmover.cfg configuration file.
User Name	Database user name. For a database on a supported Linux or UNIX system, if you have enabled PowerExchange LDAP user authentication, the user name is the enterprise user name. For more information, see the <i>PowerExchange Reference Manual</i> .
Password	<p>Password for the database user name or a valid PowerExchange passphrase.</p> <p>A PowerExchange passphrase can be from 9 to 128 characters in length and can contain the following characters:</p> <ul style="list-style-type: none"> - Uppercase and lowercase letters - The numbers 0 to 9 - Spaces - The following special characters: ' - ; # \ , . / ! % & * () _ + { } : @ < > ? <p>Note: The first character is an apostrophe.</p> <p>Passphrases cannot include single quotation marks ('), double quotation marks ("), or currency symbols.</p> <p>To use passphrases, ensure that the PowerExchange Listener runs with a security setting of SECURITY=(1,N) or higher in the DBMOVER member. For more information, see "SECURITY Statement" in the <i>PowerExchange Reference Manual</i>.</p> <p>The allowable characters in the IBM IRRPHREX exit do not affect the allowable characters in PowerExchange passphrases.</p> <p>Note: A valid RACF passphrase can be up to 100 characters in length. PowerExchange truncates passphrases longer than 100 characters when passing them to RACF for validation.</p>
Code Page	Required. Name of the code page to use for reading from or writing to the data source. Usually, this value is an ISO code page name, such as ISO-8859-6.
Pass-through security enabled	Enables pass-through security for the connection. When you enable pass-through security for a connection, the domain uses the client user name and password to log into the corresponding database, instead of the credentials defined in the connection object.
Encryption Type	<p>The type of encryption that the Data Integration Service uses. Select one of the following options:</p> <ul style="list-style-type: none"> - None - AES <p>Default is None.</p> <p>Note: Informatica recommends that you use Secure Sockets Layer (SSL) authentication instead of configuring the Encryption Type and Level connection properties. SSL authentication provides stricter security and is used by several Informatica products. For more information about implementing SSL authentication in a PowerExchange network, see the <i>PowerExchange Reference Manual</i>.</p>

Option	Description
[Encryption] Level	<p>If you select AES for Encryption Type, select one of the following options to indicate the encryption level that the Data Integration Service uses:</p> <ul style="list-style-type: none"> - 1. Use a 128-bit encryption key. - 2. Use a 192-bit encryption key. - 3. Use a 256-bit encryption key. <p>If you do not select AES for Encryption Type, this option is ignored.</p> <p>Default is 1.</p>
Pacing size	<p>Optional. Amount of data that the source system can pass to the PowerExchange Listener. Set the pacing size if an external application, database, or the Data Integration Service node is a bottleneck. User lower values for faster performance.</p> <p>The minimum value and default value is 0. A value of 0 provides the best performance.</p>
Interpret as rows	<p>Optional. Select this option to express the pacing size as a number of rows. Clear this option to express the pacing size in kilobytes. By default, this option is not selected and the pacing size is in kilobytes.</p>
Compression	<p>Optional. Select this option to enable source data compression. By compressing data, you can decrease the amount of data that Informatica applications send over the network. By default, this option is not selected and compression is disabled.</p>
Offload processing	<p>Optional. Controls whether to offload some bulk data processing from the source machine to the Data Integration Service machine. Select one of the following options:</p> <ul style="list-style-type: none"> - AUTO. The Data Integration Service determines whether to use offload processing. - Yes. Use offload processing. - No. Do not use offload processing. <p>Default is AUTO.</p>
Worker threads	<p>Optional. Number of threads that the Data Integration Service uses to process bulk data when offload processing is enabled. For optimal performance, this value should not exceed the number of available processors on the Data Integration Service machine. Valid values are 1 through 64. Default is 0, which disables multithreading.</p>
Array size	<p>Optional. The number of records in the storage array for the worker threads. This option is applicable when you set the Worker Threads option to a value greater than 0. Valid values are 1 to 5000. Default is 25.</p>
Write mode	<p>Optional. Mode in which Data Integration Service sends data to the PowerExchange Listener. Select one of the following write modes:</p> <ul style="list-style-type: none"> - CONFIRMWRITEON. Sends data to the PowerExchange Listener and waits for a response before sending more data. Select this option when error recovery is a priority. However, this option might degrade performance. - CONFIRMWRITEOFF. Sends data to the PowerExchange Listener without waiting for a response. Use this option if you can reload the target table when an error occurs. - ASYNCHRONOUSWITHFAULTTOLERANCE. Sends data to the PowerExchange Listener without waiting for a response. This option also enables error detection. This option combines the speed of CONFIRMWRITEOFF and the data integrity of CONFIRMWRITEON. <p>Default is CONFIRMWRITEON.</p>

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.
Access Key ID	Amazon S3 bucket access key ID. Note: Required if you do not use AWS Identity and Access Management (IAM) authentication.
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.
JDBC URL	Amazon Redshift connection URL.

Property	Description
Cluster Region	<p>Optional. The AWS cluster region in which the bucket you want to access resides.</p> <p>Select a cluster region if you choose to provide a custom JDBC URL that does not contain a cluster region name in the JDBC URL connection property.</p> <p>If you specify a cluster region in both Cluster Region and JDBC URL connection properties, the Data Integration Service ignores the cluster region that you specify in the JDBC URL connection property.</p> <p>To use the cluster region name that you specify in the JDBC URL connection property, select None as the cluster region in this property.</p> <p>Select one of the following cluster regions:</p> <p>Select one of the following regions:</p> <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) <p>Default is None.</p> <p>You can only read data from or write data to the cluster regions supported by AWS SDK used by PowerExchange for Amazon Redshift.</p>
Customer Master Key ID	<p>Optional. Specify the customer master key ID generated by AWS Key Management Service (AWS KMS) or the Amazon Resource Name (ARN) of your custom key for cross-account access. 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>

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.
Location	The domain where you want to create the connection.
Type	The Amazon S3 connection type.
Access Key	Access key to access the Amazon S3 bucket. Provide the access key value based on the following authentication methods: <ul style="list-style-type: none">- Basic authentication: provide the actual access key value.- IAM authentication: do not provide the access key value.- Temporary security credentials via assume role: provide access key of an IAM user with no permissions to access Amazon S3 bucket.
Secret Key	Secret access key to access the Amazon S3 bucket. The secret key is associated with the access key and uniquely identifies the account. Provide the access key value based on the following authentication methods: <ul style="list-style-type: none">- Basic authentication: provide the actual access secret value.- IAM authentication: do not provide the access secret value.- Temporary security credentials via assume role: provide access secret of an IAM user with no permissions to access Amazon S3 bucket.
IAM Role ARN	The ARN of the IAM role assumed by the user to use the dynamically generated temporary security credentials. Enter the value of this property if you want to use the temporary security credentials to access the AWS resources. If you want to use the temporary security credentials with IAM authentication, do not provide the Access Key and Secret Key connection properties. If you want to use the temporary security credentials without IAM authentication, you must enter the value of the Access Key and Secret Key connection properties. For more information about how to obtain the ARN of the IAM role, see the AWS documentation.
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.

Property	Description
S3 Account Type	<p>The type of the Amazon S3 account.</p> <p>Select Amazon S3 Storage or S3 Compatible Storage.</p> <p>Select the Amazon S3 storage option to use the Amazon S3 services. Select the S3 compatible storage option to specify the endpoint for a third-party storage provider such as Scalify RING.</p> <p>By default, Amazon S3 storage is selected.</p>
REST Endpoint	<p>The S3 storage endpoint.</p> <p>Specify the S3 storage endpoint in HTTP/HTTPS format when you select the S3 compatible storage option. For example, <code>http://s3.isv.scalify.com</code>.</p>
Region Name	<p>Select the AWS region in which the bucket you want to access resides.</p> <p>Select one of the following regions:</p> <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 (Hong Kong) - 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) <p>Default is US East (N. Virginia).</p> <p>Not applicable for S3 compatible storage.</p>
Customer Master Key ID	<p>Optional. Specify the customer master key ID or alias name generated by AWS Key Management Service (AWS KMS) or the Amazon Resource Name (ARN) of your custom key for cross-account access. You must generate the customer master key for the same region where Amazon S3 bucket reside.</p> <p>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>
Federated SSO IdP	<p>SAML 2.0-enabled identity provider for the federated user single sign-on to use with the AWS account.</p> <p>PowerExchange for Amazon S3 supports only the ADFS 3.0 identity provider.</p> <p>Select <code>None</code> if you do not want to use federated user single sign-on.</p>

Federated user single sign-on connection properties

Configure the following properties when you select `ADFS 3.0` in **Federated SSO IdP**:

Property	Description
Federated User Name	User name of the federated user to access the AWS account through the identity provider.
Federated User Password	Password for the federated user to access the AWS account through the identity provider.
IdP SSO URL	Single sign-on URL of the identity provider for AWS.
SAML Identity Provider ARN	ARN of the SAML identity provider that the AWS administrator created to register the identity provider as a trusted provider.
Role ARN	ARN of the IAM role assumed by the federated user.

Blockchain Connection Properties

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

The following table describes the general connection properties for a blockchain 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.
Swagger File Path	The absolute path of the swagger file path that contains the REST API to communicate with the blockchain. The swagger file must be a JSON file that is stored on the Data Integration Service machine. If the swagger file is in a different file format, such as YAML, convert the file to JSON format.
Base URL	Required. The base URL that is used to access assets on the blockchain.
Auth Type*	Authentication method that the run-time engine uses to connect to the REST server. You can use none, basic, digest, or OAuth.
Auth User ID*	User name to authenticate to the REST server.
Auth Password*	Password for the user name to authenticate to the REST server.

Property	Description
OAuth Consumer Key*	Required for the OAuth authentication type. Client key that is associated with the REST server.
OAuth Consumer Secret*	Required for the OAuth authentication type. Client password to connect to the REST server.
OAuth Token*	Required for the OAuth authentication type. Access token to connect to the REST server.
OAuth Token Secret*	Required for the OAuth authentication type. Password associated with the OAuth token.
Proxy Type*	Type of proxy. You can use no proxy, platform proxy, or custom.
Proxy Details*	Proxy configuration using the format <host>:<port>.
TrustStore File Path*	The absolute path of the truststore file that contains the SSL certificate.
TrustStore Password*	Password for the truststore file.
KeyStore File Path*	The absolute path of the keystore file that contains the keys and certificates required to establish a two-way secure connection with the REST server.
KeyStore Password*	Password for the keystore file.
Advanced Properties	<p>List of advanced properties to access an asset on the blockchain. Specify the advanced properties using name-value pairs that are separated by a semicolon.</p> <p>You can use the following advanced properties:</p> <ul style="list-style-type: none"> - X-API-KEY. Required if you authenticate to the REST server using an API key. <p>The advanced properties that you configure in the connection override the values for the corresponding advanced properties in the blockchain data object. For example, if the connection and the data object both specify a base URL, the value in the connection overrides the value in the data object.</p> <p>Note: The advanced properties have the precedence level, Operation level > Object level > Connection level. The properties configured at the operation level will override the properties configured at the object or connection level.</p>
Cookies	<p>Required based on how the REST API is implemented. List of cookie properties to specify the cookie information that is passed to the REST server. Specify the properties using name-value pairs that are separated by a semicolon.</p> <p>The cookie properties that you configure in the connection override the values for the corresponding cookie properties in the blockchain data object.</p>
<p>* The property is ignored. To use the functionality, configure the property as an advanced property and provide a name-value pair based on the property name in the swagger file.</p> <p>For example, configure the following name-value pair to use basic authorization:</p> <pre>Authorization=Basic <credentials></pre> <p>Note: You cannot use Test Connection to validate a Blockchain connection.</p>	

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. Not applicable for Data Engineering Streaming.
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.
SSL Mode	Select disabled . Not applicable for PowerExchange for Cassandra JDBC. SSL mode indicates the encryption type to use for the connection. You can choose a mode from the following SSL modes: <ul style="list-style-type: none">- Disabled- One way- Two way

Property	Description
SSL Truststore Path	Not applicable for PowerExchange for Cassandra JDBC or when disabled SSL mode is selected. Absolute path and file name of the SSL truststore file that contains certificates of the trusted SSL server.
SSL Truststore Password	Not applicable for PowerExchange for Cassandra JDBC or when disabled SSL mode is selected. Password for the SSL truststore.
SSL Keystore Path	Not applicable for PowerExchange for Cassandra JDBC or when disabled SSL mode is selected. Absolute path and file name of the SSL keystore file that contains private keys and certificates for the SSL server.
SSL Keystore Password	Not applicable for PowerExchange for Cassandra JDBC or when disabled SSL mode is selected. Password for the SSL keystore.
Additional Connection Properties	Enter one or more JDBC connection parameters in the following format: <pre><param1>=<value>;<param2>=<value>;<param3>=<value></pre> PowerExchange for Cassandra JDBC supports the following JDBC connection parameters: <ul style="list-style-type: none"> - BinaryColumnLength - DecimalColumnScale - EnableCaseSensitive - EnableNullInsert - EnablePaging - RowsPerPage - StringColumnLength - VTableSeparator

Confluent Kafka Connection

The Confluent Kafka connection is a Messaging connection. Use the Confluent Kafka connection to access a Kafka broker or a Confluent Kafka broker as a source or a target. You can create and manage a Confluent Kafka connection in the Developer tool or through infacmd.

When you configure a Confluent Kafka connection, you configure the following properties:

- List of Kafka brokers or Confluent Kafka brokers that the connection reads from or writes to.
- Number of seconds the Integration Service attempts to reconnect to the database if the connection fails.
- Version of the Confluent Kafka messaging broker.

General Properties

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

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	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	Description of the connection. Enter a string that you can use to identify the connection. The description cannot exceed 4,000 characters.
Location	Domain where you want to create the connection. Select the domain name.
Type	Connection type. Select <code>Messaging/ConfluentKafka</code> .

Confluent Kafka Broker Properties

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

Property	Description
Kafka Broker List	Comma-separated list of Confluent Kafka brokers that maintain the configuration of the Confluent Kafka messaging broker. To specify a Confluent Kafka broker, use the following format: <code><IP address>:<port></code>
Retry Timeout	Number of seconds after which the Data Integration Service attempts to reconnect to the Confluent Kafka broker to read or 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	Version of the Confluent Kafka messaging broker.
Additional Connection Properties	Optional. Comma-separated list of connection properties to connect to the Kafka broker.
Schema Registry URL	Location and port of the schema registry provider on which to connect.

Additional Connection Properties

You can use the following syntax for specifying the additional connection properties:

```
request.timeout.ms=<value>,session.timeout.ms=<value>,  
fetch.max.wait.ms=<value>,heartbeat.interval.ms=<value>,  
security.protocol=SASL_PLAINTEXT,sasl.kerberos.  
service.name=<kerberos name>,sasl.mechanism=GSSAPI,  
sasl.jaas.config=com.sun.security.auth.module.  
Krb5Login Modulerequired useKeyTab=true  
doNotPrompt=true storeKey=true client=true  
keyTab="<Keytab Location>" principal="<principal>;
```

SSL Properties

The following table describes the SSL properties for the Confluent Kafka connection:

Property	Description
SSL Mode	Optional. SSL mode indicates the encryption type to use for the connection. You can choose one of the following SSL modes: <ul style="list-style-type: none">- Disabled- One way- Two way The default value is <code>Disabled</code> .
SSL TrustStore File Path	Required when <code>One way</code> SSL mode is selected. Absolute path and file name of the SSL truststore file that contains certificates of the trusted SSL server.
SSL TrustStore Password	Required when <code>One way</code> SSL mode is selected. Password for the SSL truststore.
SSL KeyStore File Path	Required when <code>Two way</code> SSL mode is selected. Absolute path and file name of the SSL keystore file that contains private keys and certificates for the SSL server.
SSL KeyStore Password	Required when <code>Two way</code> SSL mode is selected. Password for the SSL keystore.
Additional Security Properties	Optional. Comma-separated list of connection properties to connect to the Confluent Kafka broker in a secured way.

Creating a Confluent Kafka Connection Using `infacmd`

You can use the `infacmd` command line program to create a Confluent Kafka connection.

To create a Confluent Kafka connection on UNIX, run the following command:

```
sh infacmd.sh createConnection -dn <domain name> -un <domain user> -pd <domain password>
-cn <connection name> -cid <connection id> -ct ConfluentKafka -o
"kfkBrkList='<host1:port1>,<host2:port2>,<host3:port3>' kafkabrokerverversion='<version>'
schemaregistryurl='<schema registry URL>'"
```

For more information about the `CreateConnection` command, see the *Informatica Command Reference*.

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.
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 property in the Databricks Spark connection.

Set to: TRUE

infaspark.pythontx.exec

Required to run a Python transformation on the Databricks Spark engine. Set to the location of the Python executable binary on the worker nodes in the Databricks cluster.

When you provision the cluster at run time, set this property in the Databricks cloud provisioning configuration. Otherwise, set on the Databricks connection.

For example, set to:

```
infaspark.pythontx.exec=/databricks/python3/bin/python3
```

infaspark.pythontx.executorEnv.PYTHONHOME

Required to run a Python transformation on the Databricks Spark engine. Set to the location of the Python installation directory on the worker nodes in the Databricks cluster.

When you provision the cluster at run time, set this property in the Databricks cloud provisioning configuration. Otherwise, set on the Databricks connection.

For example, set to:

```
infaspark.pythontx.executorEnv.PYTHONHOME=/databricks/python3
```

Greenplum Connection Properties

Use a Greenplum connection to connect to a Greenplum database. The Greenplum connection is a relational type connection. You can create and manage a Greenplum connection in the Administrator tool or the Developer tool.

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

When you create a Greenplum connection, you enter information for metadata and data access.

The following table describes Greenplum connection properties:

Property	Description
Name	Name of the Greenplum relational 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 fewer 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	Domain on which you want to create the connection.
Type	Type of connection.

The user name, password, driver name, and connection string are required to import the metadata. The following table describes the properties for metadata access:

Property	Description
User Name	User name with permissions to access the Greenplum database.
Password	Password to connect to the Greenplum database.
Driver Name	The name of the Greenplum JDBC driver. For example: <code>com.pivotal.jdbc.GreenplumDriver</code> For more information about the driver, see the Greenplum documentation.
Connection String	Use the following connection URL: <code>jdbc:pivotal:greenplum://<hostname>:<port>;DatabaseName=<database_name></code> For more information about the connection URL, see the Greenplum documentation.

PowerExchange for Greenplum uses the host name, port number, and database name to create a control file to provide load specifications to the Greenplum gpload bulk loading utility. It uses the Enable SSL option and the certificate path to establish secure communication to the Greenplum server over SSL.

The following table describes the connection properties for data access:

Property	Description
Host Name	Host name or IP address of the Greenplum server.
Port Number	Greenplum server port number. If you enter 0, the gpload utility reads from the environment variable \$PGPORT. Default is 5432.
Database Name	Name of the database.
Enable SSL	Select this option to establish secure communication between the gpload utility and the Greenplum server over SSL.
Certificate Path	Path where the SSL certificates for the Greenplum server are stored. For information about the files that need to be present in the certificates path, see the gpload documentation.

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 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 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 creates a local stage file to store the data temporarily. 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>
Dataset ID	Not applicable for PowerExchange for Google BigQuery.
Use Legacy SQL For Custom Query	Not applicable for PowerExchange for Google BigQuery.
Dataset Name for Custom Query	Not applicable for PowerExchange for Google BigQuery.
Region ID	<p>The region name where the Google BigQuery dataset resides.</p> <p>For example, if you want to connect to a Google BigQuery dataset that resides in Las Vegas region, specify us-west4 as the Region ID.</p> <p>Note: In the Storage Path connection property, ensure that you specify a bucket name or the bucket name and folder name that resides in the same region as the dataset in Google BigQuery.</p> <p>For more information about the regions supported by Google BigQuery, see the following Google BigQuery documentation:https://cloud.google.com/bigquery/docs/locations</p>

Property	Description
Optional Properties	<p>Specifies whether you can configure certain source and target functionalities through custom properties.</p> <p>You can select one of the following options:</p> <ul style="list-style-type: none"> - None. Select if you do not want to configure any custom properties. - Required. If you want to specify custom properties to configure the source and target functionalities. <p>Default is None.</p>
Provide Optional Properties	<p>Comma-separated key-value pairs of custom properties to enable certain source and target functionalities.</p> <p>Appears only when you select Required in the Optional Properties.</p> <p>For more information about the list of custom properties that you can specify, see the Informatica Knowledge Base article: https://kb.informatica.com/faq/7/Pages/26/632722.aspx</p>

Connection Modes

You can configure a Google BigQuery connection to use one of the following connection modes:

Simple mode

If you use simple mode, PowerExchange for Google BigQuery flattens each field within the Record data type field as a separate field in the Google BigQuery data object.

Hybrid mode

If you use hybrid mode, PowerExchange for Google BigQuery 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 Google BigQuery data object.

Complex mode

If you use complex mode, PowerExchange for Google BigQuery displays all the columns in the Google BigQuery table as a single field of the String data type in the Google BigQuery data object.

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	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 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.

Property	Description
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 <code>project_id</code> 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 <code>client_email</code> value present in the JSON file that you download after you create a service account.
Service Account Key	Specifies the <code>private_key</code> value present in the JSON file that you download after you create a service account.

Google PubSub Connection Properties

When you create a Google PubSub connection, you must configure the connection properties.

The following table describes the Google PubSub 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	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.
Connection Type	The connection type. Select Pub Sub connection type.
Client Email	The <code>client_email</code> value available in the JSON file that you download after you create a service account.
Client Id	The <code>client_id</code> value available in the JSON file that you download after you create a service account.
Private Key Id	The <code>private_key_id</code> value available in the JSON file that you download after you create a service account.

Property	Description
Private Key	The <code>private_key</code> value available in the JSON file that you download after you create a service account.
Project Id	The <code>project_id</code> value available 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.

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:

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

Note: When a mapping uses Hive Server 2 to run a job or parts of a job, you cannot override properties that are configured on the cluster level in preSQL or post-SQL queries or SQL override statements. Workaround: Instead of attempting to use the cluster configuration on the domain to override cluster properties, pass the override settings to the JDBC URL. For example: `beeline -u "jdbc:hive2://<domain host>:<port_number>/tpch_text_100" --hiveconf hive.execution.engine=tez`

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.
Cluster Environment Variables*	Environment variables that the Hadoop cluster uses. If you use a Cloudera CDH 6.x cluster or a Cloudera CDP cluster, configure the locale setting as cluster environment variables. In Cloudera Manager, you must also add the environment variables to the following YARN property: <code>yarn.nodemanager.env-whitelist</code> For example, the variable ORACLE_HOME represents the directory where the Oracle database client software is installed. 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: 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 Note: When a mapping uses Hive Server 2 to run a job or parts of a job, you cannot override properties that are configured on the cluster level in preSQL or post-SQL queries or SQL override statements. Workaround: Instead of attempting to use the cluster configuration on the domain to override cluster properties, pass the override settings to the JDBC URL. For example: <code>beeline -u "jdbc:hive2://<domain host>:<port_number>/tpch_text_100" --hiveconf hive.execution.engine=tez</code>
Cluster Library Path*	The path for shared libraries on the cluster. The <code>\$DEFAULT_CLUSTER_LIBRARY_PATH</code> variable contains a list of default directories.

Property	Description
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 <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: When a mapping uses Hive Server 2 to run a job or parts of a job, you cannot override properties that are configured on the cluster level in preSQL or post-SQL queries or SQL override statements. Workaround: Instead of attempting to use the cluster configuration on the domain to override cluster properties, pass the override settings to the JDBC URL. For example: <code>beeline -u "jdbc:hive2://<domain host>:<port_number>/tpch_text_100" --hiveconf hive.execution.engine=tez</code></p>
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>
<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.</p>	

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>Data Engineering Integration 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 run-time engines run mappings with the operating system profile specified Hadoop impersonation properties.</p>
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>

Property	Description
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>
Environment SQL	<p>SQL commands to set the Hadoop environment. The Data Integration Service executes the environment SQL at the beginning of each Hive script generated by a HiveServer2 job.</p> <p>The following rules and guidelines apply to the usage of environment SQL:</p> <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.
Engine Type	<p>The Data Integration Service uses HiveServer2 to process portions of some jobs by running HiveServer2 tasks on the Spark engine. When you import the cluster configuration through the admin tool, you can choose to create connections. The engine type property is populated by default based on the distribution.</p> <p>When you manually create a connection, you must configure the engine type. Use the following engine type, depending on the Hadoop distribution:</p> <ul style="list-style-type: none"> - Amazon EMR - Tez - Azure HDI - Tez - Cloudera CDH - MRv2 - Cloudera CDP - Tez - Hortonworks HDP - Tez - Mapr - MRv2
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: When a mapping uses Hive Server 2 to run a job or parts of a job, you cannot override properties that are configured on the cluster level in preSQL or post-SQL queries or SQL override statements. Workaround: Instead of attempting to use the cluster configuration on the domain to override cluster properties, pass the override settings to the JDBC URL. For example: <code>beeline -u "jdbc:hive2://<domain host>:<port_number>/tpch_text_100" --hiveconf hive.execution.engine=tez</code></p> <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.

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. Note: If YARN preemption is enabled on the cluster, verify with the Hadoop administrator that preemption is disabled on the queue associated with the Blaze engine.
Blaze Job Monitor Address	The host name and port number for the Blaze Job Monitor. Use the following format: <code><hostname>:<port></code> Where <ul style="list-style-type: none">- <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> <p>Note: You cannot use node labels on a Cloudera CDH cluster.</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: When a mapping uses Hive Server 2 to run a job or parts of a job, you cannot override properties that are configured on the cluster level in preSQL or post-SQL queries or SQL override statements. Workaround: Instead of attempting to use the cluster configuration on the domain to override cluster properties, pass the override settings to the JDBC URL. For example: <code>beeline -u "jdbc:hive2://<domain host>:<port_number>/tpch_text_100" --hiveconf hive.execution.engine=tez</code></p> <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 infacmd ms runMapping 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: When a mapping uses Hive Server 2 to run a job or parts of a job, you cannot override properties that are configured on the cluster level in preSQL or post-SQL queries or SQL override statements. Workaround: Instead of attempting to use the cluster configuration on the domain to override cluster properties, pass the override settings to the JDBC URL. For example: <code>beeline -u "jdbc:hive2://<domain host>:<port_number>/tpch_text_100" --hiveconf hive.execution.engine=tez</code></p> <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>

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	<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> <pre>~ ` ! \$ % ^ & * () - + = { [] \ : ; " ' < , > . ? /</pre>
ID	<p>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.</p>
Description	<p>The description of the connection. The description cannot exceed 4,000 characters.</p>
Location	<p>The domain where you want to create the connection.</p>

Property	Description
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.

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.
User Name	User name to access HDFS.
NameNode URI	The URI to access the storage system. You can find the value for <code>fs.defaultFS</code> in the <code>core-site.xml</code> configuration set of the cluster configuration. 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 use a Cloudera CDP Public Cloud compute cluster and the HDFS is on a Cloudera Data Lake cluster, set the property <code>spark.yarn.access.hadoopFileSystems</code> in the Spark properties of the Hadoop Connection to the same value as set here.

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	<code>hdfs://<namenode>:<port></code> where: <ul style="list-style-type: none">- <namenode> is the host name or IP address of the NameNode.- <port> is the port that the NameNode listens for remote procedure calls (RPC). <code>hdfs://<nameservice></code> in case of NameNode high availability.
MapR-FS	<code>maprfs:///</code>
WASB in HDInsight	<code>wasb://<container_name>@<account_name>.blob.core.windows.net/<path></code> where: <ul style="list-style-type: none">- <container_name> identifies a specific Azure Storage Blob container. Note: <container_name> is optional. <ul style="list-style-type: none">- <account_name> identifies the Azure Storage Blob object. Example: <code>wasb://infabdmoffering1storage.blob.core.windows.net/infabdmoffering1cluster/mr-history</code>
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 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.
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: ~ ` ! \$ % ^ & * () - + = { [] } \ : ; " ' < , > . ? /
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.
Location	The domain where you want to create the connection. Not valid for the Analyst tool.
Type	The connection type. Select Hive.
LDAP username	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. 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. If the Hadoop cluster does not use Kerberos authentication, the user name depends on the behavior of the JDBC driver. If you do not specify a user name, the Hadoop cluster authenticates jobs based on the following criteria: <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.

Property	Description
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> - 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	<p>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.</p>

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. 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>Data Engineering 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>Data Engineering 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.</p> <p>The Hive Staging Database Name is automatically updated from the Data Access Connection String. If you want to override the default name, you need to configure the Hive Staging Database Name in the Hive connection.</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>

HTTP Connection Properties

Use an HTTP connection to connect a REST Web Service Consumer transformation to a web service. The HTTP connection is a web type connection. You create an HTTP connection in the Developer tool. You can manage an HTTP connection in the Administrator tool or the Developer tool.

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

The following table describes HTTP 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.
Username	User name to connect to the web service. Enter a user name if you enable HTTP authentication or WS-Security. If the Web Service Consumer transformation includes WS-Security ports, the transformation receives a dynamic user name through an input port. The Data Integration Service overrides the user name defined in the connection.
Password	Password for the user name. Enter a password if you enable HTTP authentication or WS-Security. If the Web Service Consumer transformation includes WS-Security ports, the transformation receives a dynamic password through an input port. The Data Integration Service overrides the password defined in the connection.
End Point URL	URL for the web service that you want to access. The Data Integration Service overrides the URL defined in the WSDL file. If the Web Service Consumer transformation includes an endpoint URL port, the transformation dynamically receives the URL through an input port. The Data Integration Service overrides the URL defined in the connection.
Timeout	Number of seconds that the Data Integration Service waits for a response from the web service provider before it closes the connection. Specify a timeout value between 1 and 10,000 seconds.
HTTP Authentication Type	Type of user authentication over HTTP. Select one of the following values: <ul style="list-style-type: none"> - None. No authentication. - Automatic. The Data Integration Service chooses the authentication type of the web service provider. - Basic. Requires you to provide a user name and password for the domain of the web service provider. The Data Integration Service sends the user name and the password to the web service provider for authentication. - Digest. Requires you to provide a user name and password for the domain of the web service provider. The Data Integration Service generates an encrypted message digest from the user name and password and sends it to the web service provider. The provider generates a temporary value for the user name and password and stores it in the Active Directory on the Domain Controller. It compares the value with the message digest. If they match, the web service provider authenticates you. - NTLM. Requires you to provide a domain name, server name, or default user name and password. The web service provider authenticates you based on the domain you are connected to. It gets the user name and password from the Windows Domain Controller and compares it with the user name and password that you provide. If they match, the web service provider authenticates you. NTLM authentication does not store encrypted passwords in the Active Directory on the Domain Controller.
Trust Certificates File	File containing the bundle of trusted certificates that the Data Integration Service uses when authenticating the SSL certificate of the web service. Enter the file name and full directory path. Default is <Informatica installation directory>/services/shared/bin/ca-bundle.crt.

Property	Description
Client Certificate File Name	Client certificate that a web service uses when authenticating a client. Specify the client certificate file if the web service needs to authenticate the Data Integration Service.
Client Certificate Password	Password for the client certificate. Specify the client certificate password if the web service needs to authenticate the Data Integration Service.
Client Certificate Type	Format of the client certificate file. Select one of the following values: - PEM. Files with the .pem extension. - DER. Files with the .cer or .der extension. Specify the client certificate type if the web service needs to authenticate the Data Integration Service.
Private Key File Name	Private key file for the client certificate. Specify the private key file if the web service needs to authenticate the Data Integration Service.
Private Key Password	Password for the private key of the client certificate. Specify the private key password if the web service needs to authenticate the Data Integration Service.
Private Key Type	Type of the private key. PEM is the supported type.

IBM DB2 Connection Properties

Use an IBM DB2 connection to access IBM DB2. An IBM DB2 connection is a relational database connection. You can create and manage an IBM DB2 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 DB2 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.

Property	Description
Pass-through security enabled	Enables pass-through security for the connection. When you enable pass-through security for a connection, the domain uses the client user name and password to log into the corresponding database, instead of the credentials defined in the connection object.
Connection String for data access	<p>The DB2 connection URL used to access metadata from the database.</p> <p>dbname</p> <p>Where dbname is the alias configured in the DB2 client.</p>
Metadata Access Properties: Connection String	<p>Use the following metadata connection string URL:</p> <pre>jdbc:informatica:db2://<host name>:<port>;DatabaseName=<database name></pre> <p>When you import a table, by default, all tables are displayed under the default schema name. To view tables under a specific schema instead of the default schema, you can specify the schema name from which you want to import the table. Include the ischename parameter in the URL to specify the schema name. For example, use the following syntax to import a table from a specific schema:</p> <pre>jdbc:informatica:db2://<host name>:<port>;DatabaseName=<database name>;ischename=<schema_name></pre> <p>To search for a table in multiple schemas and import it, you can specify multiple schema names in the ischename parameter. The schema name is case sensitive. You cannot use special characters when you specify multiple schema names. Use the pipe () character to separate multiple schema names. For example, use the following syntax to search for a table in three schemas and import it:</p> <pre>jdbc:informatica:db2://<host name>:<port>;DatabaseName=<database name>;ischename=<schema_name1> <schema_name2> <schema_name3></pre> <p>When you specify multiple schema names, you must clear the Show Default Schema Only option to view the tables under the specified schema names.</p>

Property	Description
AdvancedJDBCSecurityOptions	<p>Database parameters for metadata access to a secure database. Informatica treats the value of the AdvancedJDBCSecurityOptions field as sensitive data and stores the parameter string encrypted.</p> <p>To connect to a secure database, include the following parameters:</p> <ul style="list-style-type: none"> - EncryptionMethod. Required. Indicates whether data is encrypted when transmitted over the network. This parameter must be set to SSL. - ValidateServerCertificate. Optional. Indicates whether Informatica validates the certificate that is sent by the database server. <p>If this parameter is set to True, Informatica validates the certificate that is sent by the database server. If you specify the HostNameInCertificate parameter, Informatica also validates the host name in the certificate.</p> <p>If this parameter is set to false, Informatica does not validate the certificate that is sent by the database server. Informatica ignores any truststore information that you specify.</p> <ul style="list-style-type: none"> - HostNameInCertificate. Optional. Host name of the machine that hosts the secure database. If you specify a host name, Informatica validates the host name included in the connection string against the host name in the SSL certificate. - cryptoProtocolVersion. Optional. If you enable TLS for the IBM DB2 instance, set the cryptoProtocolVersion parameter as follows: cryptoProtocolVersion=TLSv<version number>. For example, cryptoProtocolVersion=TLSv1.2 <p>Note: The version number must be the same as the TLS version you configured for the server.</p> <ul style="list-style-type: none"> - TrustStore. Required. Path and file name of the truststore file. <p>Note: If you configure SSL or TLS and specify only the file name, you must copy the truststore file to the following directory to test the connection: <Informatica server installation directory>/tomcat/bin</p> <ul style="list-style-type: none"> - TrustStorePassword. Required. Password for the truststore file for the secure database. <p>Note: Informatica appends the secure JDBC parameters to the connection string. If you include the secure JDBC parameters directly to the connection string, do not enter any parameters in the AdvancedJDBCSecurityOptions field.</p>
Data Access Properties: Connection String	<p>The connection string used to access data from the database.</p> <p>For IBM DB2 this is <database name></p>
Code Page	<p>The code page used to read from a source database or to write to a target database or file.</p>
Environment SQL	<p>SQL commands to set the database environment when you connect to the database. The Data Integration Service runs the connection environment SQL each time it connects to the database.</p>
Transaction SQL	<p>SQL commands to set the database environment when you connect to the database. The Data Integration Service runs the transaction environment SQL at the beginning of each transaction.</p>
Retry Period	<p>This property is reserved for future use.</p>
Tablespace	<p>The tablespace name of the database.</p>

Property	Description
SQL Identifier Character	Type of character that the database uses to enclose delimited identifiers in SQL queries. The available characters depend on the database type. 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. 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.
Support Mixed-case Identifiers	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. When the SQL Identifier Character property is set to none, the Support Mixed-case Identifiers property is disabled.
ODBC Provider	ODBC. The type of database to which ODBC connects. For pushdown optimization, specify the database type to enable the Data Integration Service to generate native database SQL. The options are: - Other - Sybase - Microsoft_SQL_Server Default is Other.

IBM DB2 for i5/OS Connection Properties

Use an IBM DB2 for i5/OS connection to access tables in IBM DB2 for i5/OS. An IBM DB2 for i5/OS connection is a relational database connection. You can create and manage an IBM DB2 for i5/OS connection in the Administrator tool or the Developer tool.

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

The following table describes DB2 for i5/OS 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 255 characters.
Connection Type	The connection type (DB2I).
Username	A database user name.

Property	Description
Password	<p>A password for the specified user name or a valid PowerExchange passphrase.</p> <p>A PowerExchange passphrase can be from 9 to 31 characters in length and can contain the following characters:</p> <ul style="list-style-type: none"> - Uppercase and lowercase letters - The numbers 0 to 9 - Spaces - The following special characters: ' - ; # \ , . / ! % & * () _ + { } : @ < > ? <p>Note: The first character is an apostrophe.</p> <p>Passphrases cannot include single quotation marks ('), double quotation marks ("), or currency symbols.</p> <p>To use passphrases, ensure that the PowerExchange Listener runs with a security setting of SECURITY=(1,N) or higher in the DBMOVER member. For more information, see "SECURITY Statement" in the <i>PowerExchange Reference Manual</i>.</p>
Pass-through security enabled	Enables pass-through security for the connection.
Database name	The database instance name.
Location	Node name for the location of the PowerExchange Listener that connects to DB2. The node name is defined in the first parameter of the NODE statement in the PowerExchange dbmover.cfg configuration file.
Environment SQL	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.
Database file overrides	<p>Specifies the i5/OS database file override in the following format:</p> <pre>from_file/to_library/to_file/to_member</pre> <p>Where:</p> <ul style="list-style-type: none"> - <i>from_file</i> is the file to be overridden. - <i>to_library</i> is the new library to use. - <i>to_file</i> is the file in the new library to use. - <i>to_member</i> is optional and is the member in the new library and file to use. *FIRST is used if nothing is specified. <p>You can specify up to eight unique file overrides on a single connection. A single override applies to a single source or target. When you specify more than one file override, enclose the string of file overrides in double quotes (") and include a space between each file override.</p> <p>Note: If you specify both Library List and Database File Overrides and a table exists in both, the Database File Overrides value takes precedence.</p>
Library list	<p>List of libraries that PowerExchange searches to qualify the table name for Select, Insert, Delete, or Update statements. PowerExchange searches the list if the table name is unqualified.</p> <p>Separate libraries with commas.</p> <p>Note: If you specify both Library List and Database File Overrides and a table exists in both, the Database File Overrides value takes precedence.</p>
Code Page	The code page used to read from a source database or write to a target database or file.

Property	Description
SQL Identifier character to use	Type of character that the database uses to enclose delimited identifiers in SQL queries. The available characters depend on the database type. 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. 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.
Support mixed case identifiers	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. When the SQL Identifier Character property is set to none, the Support Mixed-case Identifiers property is disabled.
Isolation level	Commit scope of the transaction. Select one of the following options: <ul style="list-style-type: none"> - None - CS. Cursor stability. - RR. Repeatable Read. - CHG. Change. - ALL Default is CS.
Encryption type	Optional. The type of encryption that the Data Integration Service uses. Select one of the following options: <ul style="list-style-type: none"> - None - AES Default is None. Note: Informatica recommends that you use Secure Sockets Layer (SSL) authentication instead of configuring the Encryption type and Encryption level connection properties. SSL authentication provides stricter security and is used by several Informatica products. For more information about implementing SSL authentication in a PowerExchange network, see the <i>PowerExchange Reference Manual</i> .
Encryption level	If you select AES for Encryption Type , select one of the following options to indicate the encryption level that the Data Integration Service uses: <ul style="list-style-type: none"> - 1. Use a 128-bit encryption key. - 2. Use a 192-bit encryption key. - 3. Use a 256-bit encryption key. If you do not select AES for Encryption Type , this option is ignored. Default is 1.
Pacing size	Optional. Amount of data that the source system can pass to the PowerExchange Listener. Set the pacing size if an external application, database, or the Data Integration Service node is a bottleneck. User lower values for faster performance. The minimum value and default value is 0. A value of 0 provides the best performance.
Interpret as rows	Optional. Select this option to express the pacing size as a number of rows. Clear this option to express the pacing size in kilobytes. By default, this option is not selected and the pacing size is in kilobytes.
Compression	Optional. Select this option to enable source data compression. By compressing data, you can decrease the amount of data that Informatica applications send over the network. By default, this option is not selected and compression is disabled.

Property	Description
Array size	Optional. The number of records in the storage array for the worker threads. This option is applicable when you set the Worker Threads option to a value greater than 0. Valid values are 25 to 5000. Default is 25.
Write mode	Optional. Mode in which the Data Integration Service sends data to the PowerExchange Listener. Select one of the following write modes: <ul style="list-style-type: none"> - CONFIRMWRITEON. Sends data to the PowerExchange Listener and waits for a response before sending more data. Select this option when error recovery is a priority. However, this option might degrade performance. - CONFIRMWRITEOFF. Sends data to the PowerExchange Listener without waiting for a response. Use this option if you can reload the target table when an error occurs. - ASYNCHRONOUSWITHFAULTTOLERANCE. Sends data to the PowerExchange Listener without waiting for a response. This option also enables error detection. This option combines the speed of CONFIRMWRITEOFF and the data integrity of CONFIRMWRITEON. Default is CONFIRMWRITEON.
Reject file	Overrides the default prefix of PWXR for the reject file. PowerExchange creates the reject file on the target machine when the write mode is ASYNCHRONOUSWITHFAULTTOLERANCE. Enter PWXDISABLE to prevent the creation of the reject files.

IBM DB2 for z/OS Connection Properties

Use an IBM DB2 for z/OS connection to access tables in IBM DB2 for z/OS. An IBM DB2 for z/OS connection is a relational database connection. You can create and manage an IBM DB2 for z/OS connection in the Administrator tool or the Developer tool.

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

The following table describes DB2 for z/OS 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	Description of the connection. The description cannot exceed 255 characters.
Connection Type	Connection type (DB2Z).
Username	Database user name.

Property	Description
Password	<p>Password for the specified user name or a valid PowerExchange passphrase.</p> <p>A PowerExchange passphrase can be from 9 to 128 characters in length and can contain the following characters:</p> <ul style="list-style-type: none"> - Uppercase and lowercase letters - The numbers 0 to 9 - Spaces - The following special characters: ' - ; # \ , . / ! % & * () _ + { } : @ < > ? <p>Note: The first character is an apostrophe.</p> <p>Passphrases cannot include single quotation marks ('), double quotation marks ("), or currency symbols.</p> <p>To use passphrases, ensure that the PowerExchange Listener runs with a security setting of SECURITY=(1,N) or higher in the DBMOVER member. For more information, see "SECURITY Statement" in the <i>PowerExchange Reference Manual</i>.</p> <p>The allowable characters in the IBM IRRPHREX exit do not affect the allowable characters in PowerExchange passphrases.</p> <p>Note: A valid RACF passphrase can be up to 100 characters in length. PowerExchange truncates passphrases longer than 100 characters when passing them to RACF for validation.</p>
Pass-through security enabled	Enables pass-through security for the connection.
DB2 Subsystem ID	Name of the DB2 subsystem.
Location	Node name for the location of the PowerExchange Listener that connects to DB2. The node name is defined in the first parameter of the NODE statement in the PowerExchange dbmover.cfg configuration file.
Environment SQL	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.
Correlation ID	Value to be concatenated to prefix PWX to form the DB2 correlation ID for DB2 requests.
Code Page	Code page used to read from a source database or write to a target database or file.
SQL identifier character to use	<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>
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>

Property	Description
Encryption type	<p>Optional. The type of encryption that the Data Integration Service uses. Select one of the following options:</p> <ul style="list-style-type: none"> - None - AES <p>Default is None.</p> <p>Note: Informatica recommends that you use Secure Sockets Layer (SSL) authentication instead of configuring the Encryption Type and Level connection properties. SSL authentication provides stricter security and is used by several Informatica products. For more information about implementing SSL authentication in a PowerExchange network, see the <i>PowerExchange Reference Manual</i>.</p>
Encryption level	<p>If you select AES for Encryption Type, select one of the following options to indicate the encryption level that the Data Integration Service uses:</p> <ul style="list-style-type: none"> - 1. Use a 128-bit encryption key. - 2. Use a 192-bit encryption key. - 3. Use a 256-bit encryption key. <p>If you do not select AES for Encryption Type, this option is ignored.</p> <p>Default is 1.</p>
Pacing size	<p>Optional. Amount of data that the source system can pass to the PowerExchange Listener. Set the pacing size if an external application, database, or the Data Integration Service node is a bottleneck. User lower values for faster performance.</p> <p>The minimum value and default value is 0. A value of 0 provides the best performance.</p>
Interpret as rows	<p>Optional. Select this option to express the pacing size as a number of rows. Clear this option to express the pacing size in kilobytes. By default, this option is not selected and the pacing size is in kilobytes.</p>
Compression	<p>Optional. Select this option to enable source data compression. By compressing data, you can decrease the amount of data that Informatica applications send over the network. By default, this option is not selected and compression is disabled.</p>
Offload processing	<p>Optional. Controls whether to offload some bulk data processing from the source machine to the Data Integration Service machine. Select one of the following options:</p> <ul style="list-style-type: none"> - AUTO. The Data Integration Service determines whether to use offload processing. - Yes. Use offload processing. - No. Do not use offload processing. <p>Default is No.</p>
Worker threads	<p>Optional. Number of threads that the Data Integration Service uses to process bulk data when offload processing is enabled. For optimal performance, this value should not exceed the number of available processors on the Data Integration Service machine. Valid values are 1 through 64. Default is 0, which disables multithreading.</p>
Array size	<p>Optional. The number of records in the storage array for the worker threads. This option is applicable when you set the Worker Threads option to a value greater than 0. Valid values are 1 to 5000. Default is 25.</p>

Property	Description
Write mode	<p>Mode in which the Data Integration Service sends data to the PowerExchange Listener. Configure one of the following write modes:</p> <ul style="list-style-type: none"> - CONFIRMWRITEON. Sends data to the PowerExchange Listener and waits for a response before sending more data. Select if error recovery is a priority. This option might decrease performance. - CONFIRMWRITEOFF. Sends data to the PowerExchange Listener without waiting for a response. Use this option when you can reload the target table if an error occurs. - ASYNCHRONOUSWITHFAULTTOLERANCE. Sends data to the PowerExchange Listener without waiting for a response. This option also provides the ability to detect errors. This provides the speed of Confirm Write Off with the data integrity of Confirm Write On. Default is CONFIRMWRITEON.
Reject file	<p>Overrides the default prefix of PWXR for the reject file. PowerExchange creates the reject file on the target machine when the write mode is ASYNCHRONOUSWITHFAULTTOLERANCE. Enter PWXDISABLE to prevent the creation of reject files.</p>

IMS Connection Properties

Use an IMS connection to access an IMS database. The IMS connection is a non-relational mainframe database type connection. The Data Integration Service connects to IMS through PowerExchange. You create an IMS connection in the Developer tool. You can manage an IMS connection in the Administrator tool or the Developer tool.

The following table describes IMS connection properties:

Option	Description
Location	Node name for the location of the PowerExchange Listener that connects to IMS. The node name is defined in the first parameter of the NODE statement in the PowerExchange dbmover.cfg configuration file.
User name	Database user name.

Option	Description
Password	<p>Password for the specified database user name or a valid PowerExchange passphrase. A PowerExchange passphrase can be from 9 to 128 characters in length and can contain the following characters:</p> <ul style="list-style-type: none"> - Uppercase and lowercase letters - The numbers 0 to 9 - Spaces - The following special characters: ' - ; # \ , . / ! % & * () _ + { } : @ < > ? <p>Note: The first character is an apostrophe.</p> <p>Passphrases cannot include single quotation marks ('), double quotation marks ("), or currency symbols.</p> <p>The allowable characters in the IBM IRRPHREX exit do not affect the allowable characters in PowerExchange passphrases.</p> <p>Note: A valid RACF passphrase can be up to 100 characters in length. PowerExchange truncates passphrases longer than 100 characters when passing them to RACF for validation.</p> <p>To use passphrases for IMS connections, ensure that the following requirements are met:</p> <ul style="list-style-type: none"> - The PowerExchange Listener must run with a security setting of SECURITY=(1,N) or higher in the DBMOVER member. For more information, see "SECURITY Statement" in the <i>PowerExchange Reference Manual</i>. - You must configure ODBA access to IMS as described in the <i>PowerExchange Navigator User Guide</i>. - You must use IMS data maps that specify IMS ODBA as the access method. Do not use data maps that specify the DL/1 BATCH access method because this access method requires the use of netport jobs, which do not support passphrases. - The IMS database must be online in the IMS control region to use ODBA access to IMS.
Code page	<p>Required. Name of the code page to use for reading from or writing to the data source. Usually, this value is an ISO code page name, such as ISO-8859-6.</p>
Pass-through security enabled	<p>Enables pass-through security for the connection.</p>
Encryption type	<p>The type of encryption that the Data Integration Service uses. Select one of the following options:</p> <ul style="list-style-type: none"> - None - AES <p>Default is None.</p> <p>Note: Informatica recommends that you use Secure Sockets Layer (SSL) authentication instead of configuring the Encryption Type and Level connection properties. SSL authentication provides stricter security and is used by several Informatica products. For more information about implementing SSL authentication in a PowerExchange network, see the <i>PowerExchange Reference Manual</i>.</p>
[Encryption] Level	<p>If you select AES for Encryption Type, select one of the following options to indicate the encryption level that the Data Integration Service uses:</p> <ul style="list-style-type: none"> - 1. Use a 128-bit encryption key. - 2. Use a 192-bit encryption key. - 3. Use a 256-bit encryption key. <p>If you do not select AES for Encryption Type, this option is ignored.</p> <p>Default is 1.</p>

Option	Description
Pacing size	Optional. Amount of data that the source system can pass to the PowerExchange Listener. Set the pacing size if an external application, database, or the Data Integration Service node is a bottleneck. User lower values for faster performance. The minimum value and default value is 0. A value of 0 provides the best performance.
Interpret as rows	Optional. Select this option to express the pacing size as a number of rows. Clear this option to express the pacing size in kilobytes. By default, this option is not selected and the pacing size is in kilobytes.
Compression	Optional. Select this option to enable source data compression. By compressing data, you can decrease the amount of data that Informatica applications send over the network. By default, this option is not selected and compression is disabled.
Offload processing	Optional. Controls whether to offload some bulk data processing from the source machine to the Data Integration Service machine. Select one of the following options: <ul style="list-style-type: none"> - AUTO. The Data Integration Service determines whether to use offload processing. - Yes. Use offload processing. - No. Do not use offload processing. Default is AUTO.
Worker threads	Optional. Number of threads that the Data Integration Service uses to process bulk data when offload processing is enabled. For optimal performance, this value should not exceed the number of available processors on the Data Integration Service machine. Valid values are 1 through 64. Default is 0, which disables multithreading.
Array size	Optional. The number of records in the storage array for the worker threads. This option is applicable when you set the Worker Threads option to a value greater than 0. Valid values are 1 to 5000. Default is 25.
Write mode	Optional. Mode in which Data Integration Service sends data to the PowerExchange Listener. Select one of the following write modes: <ul style="list-style-type: none"> - CONFIRMWRITEON. Sends data to the PowerExchange Listener and waits for a response before sending more data. Select this option when error recovery is a priority. However, this option might degrade performance. - CONFIRMWRITEOFF. Sends data to the PowerExchange Listener without waiting for a response. Use this option if you can reload the target table when an error occurs. - ASYNCHRONOUSWITHFAULTTOLERANCE. Sends data to the PowerExchange Listener without waiting for a response. This option also enables error detection. This option combines the speed of CONFIRMWRITEOFF and the data integrity of CONFIRMWRITEON. Default is CONFIRMWRITEON.

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	Name of the JDBC driver class. The following list provides the driver class name that you can enter for the applicable database type: <ul style="list-style-type: none"> - DataDirect JDBC driver class name for Oracle: com.informatica.jdbc.oracle.OracleDriver - DataDirect JDBC driver class name for IBM DB2: com.informatica.jdbc.db2.DB2Driver - DataDirect JDBC driver class name for Microsoft SQL Server: com.informatica.jdbc.sqlserver.SQLServerDriver - DataDirect JDBC driver class name for Sybase ASE: com.informatica.jdbc.sybase.SybaseDriver - DataDirect JDBC driver class name for Informix: com.informatica.jdbc.informix.InformixDriver - DataDirect JDBC driver class name for MySQL: com.informatica.jdbc.mysql.MySQLDriver - JDBC driver for Databricks Delta Lake: the name of the driver that you downloaded from Databricks. For information about the driver, see the topic on configuring storage access in the "Before You Begin Databricks Integration" chapter of the <i>Data Engineering Integration Guide</i>. For more information about which driver class to use with specific databases, see the vendor documentation.
Connection String	Connection string to connect to the database. Use the following connection string: jdbc:<subprotocol>:<subname> For more information about the connection string to use with specific drivers, see the vendor documentation.
Environment SQL	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. Note: If you enable Sqoop, Sqoop ignores this property.
Transaction SQL	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. Note: If you enable Sqoop, Sqoop ignores this property.

Property	Description
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>
Use Sqoop Connector	<p>Enables Sqoop connectivity for the data object that uses the JDBC connection. The Data Integration Service runs the mapping in the Hadoop run-time environment through Sqoop.</p> <p>You can configure Sqoop connectivity for relational data objects, customized data objects, and logical data objects that are based on a JDBC-compliant database.</p> <p>Select Sqoop v1.x to enable Sqoop connectivity.</p> <p>Default is None.</p>
Sqoop Arguments	<p>Enter the arguments that Sqoop must use to connect to the database. Separate multiple arguments with a space.</p> <p>To run the mapping on the Blaze engine with the Teradata Connector for Hadoop (TDCH) specialized connectors for Sqoop, you must define the TDCH connection factory class in the Sqoop arguments. The connection factory class varies based on the TDCH Sqoop Connector that you want to use.</p> <ul style="list-style-type: none"> - To use Cloudera Connector Powered by Teradata, configure the following Sqoop argument: <ul style="list-style-type: none"> - <code>Dsqoop.connection.factories=com.cloudera.connector.teradata.TeradataManagerFactory</code> - To use Hortonworks Connector for Teradata (powered by the Teradata Connector for Hadoop), configure the following Sqoop argument: <ul style="list-style-type: none"> - <code>Dsqoop.connection.factories=org.apache.sqoop.teradata.TeradataManagerFactory</code> <p>To run the mapping on the Spark engine, you do not need to define the TDCH connection factory class in the Sqoop arguments. The Data Integration Service invokes the Cloudera Connector Powered by Teradata and Hortonworks Connector for Teradata (powered by the Teradata Connector for Hadoop) by default.</p> <p>Note: To run the mapping with a generic JDBC connector instead of the specialized Cloudera or Hortonworks connector, you must define the <code>--driver</code> and <code>--connection-manager</code> Sqoop arguments in the JDBC connection. If you define the <code>--driver</code> and <code>--connection-manager</code> arguments in the Read or Write transformation of the mapping, Sqoop ignores the arguments.</p> <p>If you do not enter Sqoop arguments, the Data Integration Service constructs the Sqoop command based on the JDBC connection properties.</p>

JDBC V2 Connection Properties

When you set up a JDBC V2 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 JDBC V2 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 JDBC V2.

The **Details** tab contains the connection attributes of the JDBC V2 connection. The following table describes the connection attributes:

Property	Description
Username	The database user name. User name with permissions to access the database that supports the Type 4 JDBC driver.
Password	The password for the database user name.
Schema Name	Optional. The schema name to connect in the database. If you do not specify the schema name, all the schemas available in the database are listed.
JDBC Driver Class Name	Name of the JDBC driver class. The following list provides the driver class name that you can enter for the applicable database type: <ul style="list-style-type: none">- JDBC driver class name for Azure SQL Database: <code>com.microsoft.sqlserver.jdbc.SQLServerDriver</code>- JDBC driver class name for Aurora PostgreSQL: <code>org.postgresql.Driver</code>- JDBC driver class name for SAP HANA Database: <code>com.sap.db.jdbc.Driver</code> For more information about which driver class to use with specific databases, see the third-party vendor documentation.

Property	Description
Connection String	<p>Connection string to connect to the database.</p> <p>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 Azure SQL Database JDBC driver: <pre>jdbc:sqlserver://<host>:<port>;database=<database_name></pre> - Connection string for Aurora PostgreSQL JDBC driver: <pre>jdbc:postgresql://<host>:<port>[/<database_name>]</pre> - Connection string for SAP HANA Database driver: <pre>jdbc:sap://<host>:<port>/?databaseName=<Database_Name></pre> <p>For more information about the connection string to use with specific drivers, see the third-party vendor documentation.</p>
Sub Type	<p>The database type to which you want to connect.</p> <p>You can select from the following database types to connect:</p> <ul style="list-style-type: none"> - Azure SQL Database. Connects to Azure SQL database. - PostgreSQL. Connects to Aurora PostgreSQL database. - SAP HANA Database. Connects to SAP HANA database. - Others . Connects to any database that supports the Type 4 JDBC driver.
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>For example, Aurora PostgreSQL database supports mixed-cased characters. You must enable this property to connect to the Aurora PostgreSQL database.</p> <p>When the SQL Identifier Character property is set to none, the Support Mixed-case Identifiers property is disabled.</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: Select SQL Identifier Character as None when you specify the SAP HANA Database subtype.</p>

JD Edwards EnterpriseOne Connection Properties

Use a JD Edwards EnterpriseOne connection to connect to a JD Edwards EnterpriseOne object.

The following table describes the JD Edwards EnterpriseOne connection properties:

Property	Description
Name	The name of the connection. The name is not case sensitive and must be unique within the domain. It 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. The description cannot exceed 765 characters.
Location	The Informatica domain where you want to create the connection.
Type	The connection type. Select JD Edwards EnterpriseOne.
Host Name	JD Edwards EnterpriseOne server host name.
Enterprise Port	JD Edwards EnterpriseOne server port number. Default is 6016.
User Name	The JD Edwards EnterpriseOne database user name.
Password	The password for the JD Edwards EnterpriseOne database user.
Environment	Name of the JD Edwards EnterpriseOne environment you want to connect to.
Role	Role of the JD Edwards EnterpriseOne user. Default is *ALL.
User Name	The JD Edwards EnterpriseOne database user name.
Password	Password for the database user.

Property	Description
Driver Class Name	<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> <p>For more information about which driver class to use with specific databases, see the vendor documentation.</p>
Connection String	<p>The connection string to connect to the database. Use the following connection string:</p> <p>The JDBC connection string uses the following syntax:</p> <ul style="list-style-type: none"> - For Oracle: <code>jdbc:informatica:oracle://<host name>:<port>,ServiceName=<db service name></code> - For DB2: <code>jdbc:informatica:db2://<host name>:<port>;databaseName=<db name></code> - For Microsoft SQL: <code>jdbc:informatica:sqlserver://<host name>:<port>;databaseName=<db name></code>

Kafka Connection Properties

The Kafka connection is a Messaging connection. Use the Kafka connection to access an Apache Kafka broker as a source or a target. You can create and manage a Kafka connection in the Developer tool or through infacmd.

When you configure a Kafka connection, you configure the following properties:

- The list of Kafka brokers that the connection reads from or writes to.
- The number of seconds the Integration Service attempts to reconnect to the database if the connection fails.
- Version of the Kafka messaging broker. Configure the Kafka messaging broker version to Apache 0.10.1.1 & above.

General Properties

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. Select the domain name.
Type	The connection type. Select Messaging/Kafka.

Kafka Broker Properties

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

Property	Description
Kafka Broker List	Comma-separated list of Kafka brokers which maintain the configuration of the Kafka messaging broker. To specify a Kafka broker, use the following format: <IP Address>:<port>
Retry Timeout	Number of seconds after which the Integration Service attempts to reconnect to the Kafka broker to read or write data. If the source or target is not available for the time you specify, the mapping execution stops to avoid any data loss.

Property	Description
Kafka Broker Version	Configure the Kafka messaging broker version to Apache 0.10.1.1 & above.
Additional Connection Properties	<p>Optional. Comma-separated list of connection properties to connect to the Kafka broker.</p> <p>For example, you can use the following syntax:</p> <pre>request.timeout.ms=<value>,session.timeout.ms=<value>, fetch.max.wait.ms=<value>,heartbeat.interval.ms=<value>, security.protocol=SASL_PLAINTEXT,sasl.kerberos.service.name=<kerberos_name>,sasl.mechanism=GSSAPI, sasl.jaas.config=com.sun.security.auth.module.Krb5LoginModule;required useKeyTab=true doNotPrompt=true storeKey=true client=true keyTab="<Keytab Location>" principal="<principal>";</pre> <p>To reduce the time taken to connect to the Kafka broker, ensure that you set the following properties:</p> <ul style="list-style-type: none"> - request.timeout.ms - session.timeout.ms - fetch.max.wait.ms - heartbeat.interval.ms <p>To connect to the Kafka broker in a secured way, ensure that you set one of the following values for the <code>security.protocol</code> property:</p> <ul style="list-style-type: none"> - SASL_SSL - SSL <p>The default value of <code>security.protocol</code> property is SASL_PLAINTEXT.</p> <p>Technical Preview: The Additional Connection Properties is available for technical preview. Technical preview functionality is supported but is unwarranted and is not production-ready. Informatica recommends that you use in non-production environments only.</p> <p>For more information about the connection properties, see https://kafka.apache.org/documentation/.</p>

SSL Properties

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

Property	Description
SSL Mode	<p>Required. SSL mode indicates the encryption type to use for the connection.</p> <p>You can choose a mode from the following SSL modes:</p> <ul style="list-style-type: none"> - Disabled - One way - Two way
SSL TrustStore File Path	<p>Required when One way SSL mode is selected.</p> <p>Absolute path and file name of the SSL truststore file that contains certificates of the trusted SSL server.</p>
SSL TrustStore Password	<p>Required when One way SSL mode is selected.</p> <p>Password for the SSL truststore.</p>

Property	Description
SSL KeyStore File Path	Required when Two way SSL mode is selected. Absolute path and file name of the SSL keystore file that contains private keys and certificates for the SSL server.
SSL KeyStore Password	Required when Two way SSL mode is selected. Password for the SSL keystore.

Creating a Kafka Connection Using infacmd

You can use the infacmd command line program to create a Kafka connection.

To create a Kafka connection on UNIX, run the following command:

```
sh infacmd.sh createConnection -dn <domain name> -un <domain user> -pd <domain password> -cn
<connection name> -cid <connection id> -ct Kafka -o
kfkBrkList=<host1:port1>,<host2:port2>,<host3:port3> kafkabrokerverversion=<version>
additionalConnectionProperties=<additional properties>
```

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

Kudu Connection Properties

Use a Kudu connection to access Kudu.

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

You can create and manage a Kudu connection in the Administrator tool or the Developer tool. The following table describes the Kudu 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 Kudu.

The following table describes the properties for metadata access:

Property	Description
Kudu Master URLs	The URLs of the Kudu master tables.
Kudu Library Version	The version number of the Kudu library.
Cluster Configuration	The Hadoop cluster that you use for the connection.

LDAP Connection Properties

Use an LDAP connection to connect to an LDAP object.

The following table describes the LDAP connection properties:

Property	Description
Name	The name of the connection. The name is not case sensitive and must be unique within the domain. It 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. The description cannot exceed 765 characters.
Location	The Informatica domain where you want to create the connection.
Type	The connection type. Select LDAP.
Host Name	LDAP directory server host name. Default is localhost.
Port	LDAP directory server port number. Default is 389.
Anonymous Connection	Establishes an anonymous connection with the LDAP directory server. Select anonymous connection to access a directory server as an anonymous user without authentication. Note: You cannot establish an anonymous connection with Active Directory.
User Name	The LDAP user name to connect to the LDAP directory server.
Password	The password to connect to the LDAP directory server.
Secure Connection	Establishes a secure connection with the LDAP directory server through the TLS protocol.

Property	Description
TrustStore File Name	The file name of the truststore that contains the TLS certificate to establish a secure connection with the LDAP directory server. Default is <code>infa_truststore.jks</code> . Required if you select Secure Connection. Contact the LDAP Administrator for the truststore file name and password.
TrustStore Password	The password for the truststore file that contains the SSL certificate.
KeyStore File Name	The file name of the keystore that contains the keys and certificates required to establish a secure communication with the LDAP directory server. Required if you select Secure Connection. Contact the LDAP Administrator for the keystore file name and password.
KeyStore Password	The password for the keystore file required for secure communication.

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 Azure Blob Storage.

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.
Authorization Type	Authorization type. You can select any of the following authorization mechanisms: <ul style="list-style-type: none"> - Shared Key Authorization - Shared Access Signatures
Account Key	Microsoft Azure Storage access key. Applies when you select shared key authorization.
SAS Token	SAS URI with SAS token that you generate on Microsoft Azure portal for your account. Applies when you select shared access signature authorization type. <p>Note: You must provide a valid SAS URI with a valid SAS token.</p>
Container Name	The root container or sub-folders with the absolute path. <p>Note: To import complex files, specify only the root container.</p>
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.

Property	Description
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 Storage Gen1 Connection Properties

Use a Microsoft Azure Data Lake Storage Gen1 connection to access a Microsoft Azure Data Lake Storage Gen1.

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 Data Lake Storage Gen1 connection in the Administrator tool or the Developer tool. The following table describes the Microsoft Azure Data Lake Storage Gen1 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 Storage Gen1.

The following table describes the properties for metadata access:

Property	Description
ADLS Account Name	The name of the Microsoft Azure Data Lake Storage Gen1.
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.

Property	Description
Directory	The Microsoft Azure Data Lake Storage Gen1 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 Storage Gen1 documentation.

Microsoft Azure Data Lake Storage Gen2 Connection Properties

Use a Microsoft Azure Data Lake Storage Gen2 connection to access a Microsoft Azure Data Lake Storage Gen2.

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 Data Lake Storage Gen2 connection in the Administrator tool or the Developer tool. The following table describes the Microsoft Azure Data Lake Storage Gen2 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 Storage Gen2.

The following table describes the properties for metadata access:

Property	Description
Account Name	The Microsoft Azure Data Lake Storage Gen2 account name or the service name.
Client ID	The ID of your application to complete the OAuth Authentication in the Azure Active Directory (AD).

Property	Description
Client Secret	The client secret key to complete the OAuth Authentication in the Azure AD.
Tenant ID	The Directory ID of the Azure AD.
File System Name	The name of an existing file system in the Microsoft Azure Data Lake Storage Gen2.
Directory Path	The path of an existing directory without the file system name. There is no default directory. You can select one of the following syntax: <ul style="list-style-type: none"> - / for root directory. - /dir1 - dir1/dir2
Adls Gen2 End-point	Type of Microsoft Azure endpoints. You can select any of the following endpoints: <ul style="list-style-type: none"> - core.windows.net: Default - core.usgovcloudapi.net: To select the Azure Government endpoints

For more information about creating a client ID, client secret, tenant ID, and file system name, contact the Azure administrator or see Microsoft Azure Data Lake Storage Gen2 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 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: <code>jdbc:sqlserver://<Server>.database.windows.net:1433;database=<Database></code> . 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 Storage Type	Type of Azure storage to stage the files. You can select any of the following storage type: <ul style="list-style-type: none"> - Azure Blob. Default. To use Microsoft Azure Blob Storage to stage the files. - ADLS Gen2. To use Microsoft Azure Data Lake Storage Gen2 as storage to stage the files.
Azure Blob Account Name	Name of the Microsoft Azure Storage account to stage the files.
Azure Blob Account Key	The key that authenticates the access to the Blob storage account.
ADLS Gen2 Storage Account Name	Name of the Microsoft Azure Data Lake Storage Gen2 account to stage the files.
ADLS Gen2 Account Key	Microsoft Azure Data Lake Storage Gen2 access key to stage the files.
Blob End-point	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 You can configure the US government Microsoft Azure end-points when a mapping runs in the native environment and on the Spark engine.
VNet Rule	Enable to connect to a Microsoft Azure SQL Data Warehouse endpoint residing in a virtual network (VNet).

MS SQL Server Connection Properties

Use a Microsoft SQL Server connection to access Microsoft SQL Server. A Microsoft SQL Server connection is a connection to a Microsoft SQL Server relational database. You can create and manage a Microsoft SQL Server connection in the Administrator tool or the Developer tool.

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

The following table describes MS SQL Server 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.
Use trusted connection	Enables the application service to use Windows authentication to access the database. The user name that starts the application service must be a valid Windows user with access to the database. By default, this option is cleared. Note: Windows and NTLM authentication is not certified for a Microsoft SQL Server 2017 version hosted on Linux.
User Name	The database user name. Required if Microsoft SQL Server uses NTLMv1 or NTLMv2 authentication.
Password	The password for the database user name. Required if Microsoft SQL Server uses NTLMv1 or NTLMv2 authentication.
Pass-through security enabled	Enables pass-through security for the connection. When you enable pass-through security for a connection, the domain uses the client user name and password to log into the corresponding database, instead of the credentials defined in the connection object.
Metadata Access Properties: Connection String	<p>Connection string used to access metadata from the database.</p> <p>Use the following connection string:</p> <pre>jdbc:informatica:sqlserver://<host name>:<port>;DatabaseName=<database name></pre> <p>To test the connection with NTLM authentication, include the following parameters in the connection string:</p> <ul style="list-style-type: none"> - AuthenticationMethod. The NTLM authentication version to use. Note: UNIX supports NTLMv1 and NTLMv2 but not NTLM. - Domain. The domain that the SQL server belongs to. <p>The following example shows the connection string for a SQL server that uses NTLMv2 authentication in a NT domain named Informatica.com:</p> <pre>jdbc:informatica:sqlserver://host01:1433;DatabaseName=SQL1;AuthenticationMethod=ntlm2java;Domain=Informatica.com</pre> <p>If you connect with NTLM authentication, you can enable the Use trusted connection option in the MS SQL Server connection properties. If you connect with NTLMv1 or NTLMv2 authentication, you must provide the user name and password in the connection properties.</p>

Property	Description
AdvancedJDBCSecurityOptions	<p>Database parameters for metadata access to a secure database. Informatica treats the value of the AdvancedJDBCSecurityOptions field as sensitive data and stores the parameter string encrypted.</p> <p>To connect to a secure database, include the following parameters:</p> <ul style="list-style-type: none"> - EncryptionMethod. Required. Indicates whether data is encrypted when transmitted over the network. This parameter must be set to SSL. - ValidateServerCertificate. Optional. Indicates whether Informatica validates the certificate that is sent by the database server. <p>If this parameter is set to True, Informatica validates the certificate that is sent by the database server. If you specify the HostNameInCertificate parameter, Informatica also validates the host name in the certificate.</p> <p>If this parameter is set to false, Informatica does not validate the certificate that is sent by the database server. Informatica ignores any truststore information that you specify.</p> <ul style="list-style-type: none"> - HostNameInCertificate. Optional. Host name of the machine that hosts the secure database. If you specify a host name, Informatica validates the host name included in the connection string against the host name in the SSL certificate. - cryptoProtocolVersion. Optional. If you enable TLS for the Microsoft SQL Server instance, set the cryptoProtocolVersion parameter as follows: cryptoProtocolVersion=TLsv<version number>. For example, cryptoProtocolVersion=TLsv1.2 <p>Note: The version number must be the same as the TLS version you configured for the server.</p> <ul style="list-style-type: none"> - TrustStore. Required. Path and file name of the truststore file. <p>Note: If you configure SSL or TLS and specify only the file name, you must copy the truststore file to the following directory to test the connection: <Informatica server installation directory>/tomcat/bin</p> <ul style="list-style-type: none"> - TrustStorePassword. Required. Password for the truststore file for the secure database. <p>Not applicable for ODBC.</p> <p>Note: Informatica appends the secure JDBC parameters to the connection string. If you include the secure JDBC parameters directly to the connection string, do not enter any parameters in the AdvancedJDBCSecurityOptions field.</p>
Data Access Properties: Provider Type	<p>The connection provider that you want to use to connect to the Microsoft SQL Server database.</p> <p>You can select the following provider types:</p> <ul style="list-style-type: none"> - ODBC - Oldeb(Deprecated) <p>Default is ODBC.</p> <p>Note: Although the Microsoft SQL Server connection user interface shows the OLEDB provider type as deprecated, Informatica supports the OLEDB provider type. For more information about the OLEDB provider type support statement, see the following Knowledge Base article KB 522895.</p>
Use DSN	<p>Enables the Data Integration Service to use the Data Source Name for the connection.</p> <p>If you select the Use DSN option, the Data Integration Service retrieves the database and server names from the DSN.</p> <p>If you do not select the Use DSN option, you must provide the database and server names.</p>

Property	Description
Connection String	<p>Use the following connection string if you do not enable DSN mode:</p> <pre><server name>@<database name></pre> <p>If you enable DSN mode, use the following connection strings:</p> <pre><DSN Name></pre>
Code Page	The code page used to read from a source database or to write to a target database or file.
Domain Name	The name of the domain.
Packet Size	The packet size used to transmit data. Used to optimize the native drivers for Microsoft SQL Server.
Owner Name	<p>The name of the owner of the schema.</p> <p>Note: When you generate a table DDL through a dynamic mapping or through the Generate and Execute DDL option, the DDL metadata does not include schema name and owner name properties.</p>
Schema Name	<p>The name of the schema in the database. You must specify the schema name for the Profiling Warehouse if the schema name is different from the database user name. You must specify the schema name for the data object cache database if the schema name is different from the database user name and if you configure user-managed cache tables.</p> <p>Note: When you generate a table DDL through a dynamic mapping or through the Generate and Execute DDL option, the DDL metadata does not include schema name and owner name properties.</p>
Environment SQL	SQL commands to set the database environment when you connect to the database. The Data Integration Service runs the connection environment SQL each time it connects to the database.
Transaction SQL	SQL commands to set the database environment when you connect to the database. The Data Integration Service runs the transaction environment SQL at the beginning of each transaction.
Retry Period	This property is reserved for future use.
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>

Property	Description
Support Mixed-case Identifiers	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. When the SQL Identifier Character property is set to none, the Support Mixed-case Identifiers property is disabled.
ODBC Provider	ODBC. The type of database to which ODBC connects. For pushdown optimization, specify the database type to enable the Data Integration Service to generate native database SQL. The options are: <ul style="list-style-type: none"> - Other - Sybase - Microsoft_SQL_Server Default is Other.

Netezza Connection Properties

Use a Netezza connection to access a Netezza database. The Netezza connection is a database connection. You can create and manage a Netezza connection in the Administrator tool or the Developer tool.

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

The following table describes the Netezza 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	Description of the connection. The description cannot exceed 765 characters.
Location	Domain where you want to create the connection.
Type	Connection type. Select Netezza .
User name	User name with the appropriate permissions to access the Netezza database.
Password	Password for the database user name.
JDBC Url	JDBC URL that the Developer tool must use when it connects to the Netezza database. Use the following format: jdbc:netezza://<hostname>:<port>/<database name>

Property	Description
Connection String	Name of the ODBC data source that you want to use to connect to the Netezza database.
Timeout	Number of seconds that the Developer tool waits for a response from the Netezza database before it closes the connection.

OData Connection Properties

Use an OData connection to access an OData URL. The OData connection is a Web connection. You can create and manage an OData connection in the Administrator tool or the Developer tool.

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

The following table describes the OData 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	Description of the connection. The description cannot exceed 4,000 characters.
Location	Domain where you want to create the connection.
Type	Connection type. Select OData .
User name	Optional. User name with the appropriate permissions to read data from the OData resource.
Password	Optional. Password for the OData URL user name.
URL	OData service root URL that exposes the data that you want to read.
Security Type	Optional. Security protocol that the Developer tool must use to establish a secure connection with the OData server. Select one of the following values: - None - SSL - TLS Default is None.
TrustStore File Name	Required if you select a security type. Name of the truststore file that contains the public certificate for the OData server. Default is <code>infa_truststore.jks</code> .

Property	Description
Password	Required if you select a security type. Password for the truststore file that contains the public certificate for the OData server.
KeyStore File Name	Required if you select a security type. Name of the keystore file that contains the private key for the OData server. Default is <code>infa_truststore.jks</code> .
Password	Required if you select a security type. Password for the keystore file that contains the private key for the OData server.

ODBC Connection Properties

Use an ODBC connection to access ODBC data. An ODBC connection is a relational database connection. You can create and manage an ODBC 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 ODBC 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.
Pass-through security enabled	Enables pass-through security for the connection. When you enable pass-through security for a connection, the domain uses the client user name and password to log into the corresponding database, instead of the credentials defined in the connection object.
Data Access Properties: Connection String	The ODBC connection URL used to access metadata from the database. <data source name>
Code Page	The code page used to read from a source database or to write to a target database or file.

Property	Description
Environment SQL	SQL commands to set the database environment when you connect to the database. The Data Integration Service runs the connection environment SQL each time it connects to the database.
Transaction SQL	SQL commands to set the database environment when you connect to the database. The Data Integration Service runs the transaction environment SQL at the beginning of each transaction.
Retry Period	This property is reserved for future use.
SQL Identifier Character	Type of character that the database uses to enclose delimited identifiers in SQL queries. The available characters depend on the database type. 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. 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.
Support Mixed-case Identifiers	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. When the SQL Identifier Character property is set to none, the Support Mixed-case Identifiers property is disabled.
ODBC Provider	The type of database to which ODBC connects. For pushdown optimization, specify the database type to enable the Data Integration Service to generate native database SQL. The options are: <ul style="list-style-type: none"> - Other - Sybase - Microsoft_SQL_Server - Snowflake Default is Other.

Note: Use an ODBC connection to connect to Microsoft SQL Server when the Data Integration Service runs on UNIX or Linux. Use a native connection to Microsoft SQL Server when the Data Integration Service runs on Windows.

Oracle Connection Properties

Use an Oracle connection to connect to an Oracle database. The Oracle connection is a relational connection type. You can create and manage an Oracle 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 Oracle 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.
Pass-through security enabled	Enables pass-through security for the connection. When you enable pass-through security for a connection, the domain uses the client user name and password to log into the corresponding database, instead of the credentials defined in the connection object.
Metadata Access Properties: Connection String	<p>Connection string used to access metadata from the database.</p> <p>Use the following connection string:</p> <pre>jdbc:informatica:oracle://<host_name>:<port>;SID=<database name></pre> <p>Use the following connection string to connect to Oracle database through Oracle Connection Manager:</p> <pre>jdbc:informatica:oracle:TNSNamesFile=<fully qualified path to the tnsnames.ora file>;TNSServerName=<TNS server name>;</pre>

Property	Description
AdvancedJDBCSecurityOptions	<p>Database parameters for metadata access to a secure database. Informatica treats the value of the AdvancedJDBCSecurityOptions field as sensitive data and stores the parameter string encrypted.</p> <p>To connect to a secure database, include the following parameters:</p> <ul style="list-style-type: none"> - EncryptionMethod. Required. Indicates whether data is encrypted when transmitted over the network. This parameter must be set to SSL. - ValidateServerCertificate. Optional. Indicates whether Informatica validates the certificate that is sent by the database server. <p>If this parameter is set to True, Informatica validates the certificate that is sent by the database server. If you specify the HostNameInCertificate parameter, Informatica also validates the host name in the certificate.</p> <p>If this parameter is set to false, Informatica does not validate the certificate that is sent by the database server. Informatica ignores any truststore information that you specify.</p> <ul style="list-style-type: none"> - HostNameInCertificate. Optional. Host name of the machine that hosts the secure database. If you specify a host name, Informatica validates the host name included in the connection string against the host name in the SSL certificate. - cryptoProtocolVersion. Optional. If you enable TLS for the Oracle instance, set the cryptoProtocolVersion parameter as follows: cryptoProtocolVersion=TLSv<version number>. For example, cryptoProtocolVersion=TLSv1.2 <p>Note: The version number must be the same as the TLS version you configured for the server.</p> <ul style="list-style-type: none"> - TrustStore. Required. Path and file name of the truststore file. <p>Note: If you configure SSL or TLS and specify only the file name, you must copy the truststore file to the following directory to test the connection: <Informatica server installation directory>/tomcat/bin</p> <ul style="list-style-type: none"> - TrustStorePassword. Required. Password for the truststore file for the secure database. - KeyStore. Required. Path and file name of the keystore file. - KeyStorePassword. Required. Password for the keystore file for the secure database. <p>Note: Informatica appends the secure JDBC parameters to the connection string. If you include the secure JDBC parameters directly to the connection string, do not enter any parameters in the AdvancedJDBCSecurityOptions field.</p>
Data Access Properties: Connection String	<p>Use the following connection string:</p> <pre><database name>.world</pre>
Code Page	<p>The code page used to read from a source database or to write to a target database or file.</p>
Environment SQL	<p>SQL commands to set the database environment when you connect to the database. The Data Integration Service runs the connection environment SQL each time it connects to the database.</p>
Transaction SQL	<p>SQL commands to set the database environment when you connect to the database. The Data Integration Service runs the transaction environment SQL at the beginning of each transaction.</p>
Retry Period	<p>This property is reserved for future use.</p>

Property	Description
Enable Parallel Mode	Enables parallel processing when loading data into a table in bulk mode. By default, this option is cleared.
SQL Identifier Character	Type of character that the database uses to enclose delimited identifiers in SQL queries. The available characters depend on the database type. 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. 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.
Support Mixed-case Identifiers	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. When the SQL Identifier Character property is set to none, the Support Mixed-case Identifiers property is disabled.

Salesforce Connection Properties

Use a Salesforce connection to connect to a Salesforce object. The Salesforce connection is an application connection type. You can create and manage a Salesforce connection in the Administrator tool or the Developer tool.

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

The following table describes Salesforce connection properties:

Property	Description
Name	Name of the connection. The name is not case sensitive and must be unique within the domain. It 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 Informatica domain where you want to create the connection.
Type	The connection type. You can select the Standard connection type or OAuth connection type.
User Name	Applicable for Standard connection type. Salesforce user name.

Property	Description
User Password	Applicable for Standard connection type. Password for the Salesforce user name. To access Salesforce outside the trusted network of your organization, you must append a security token to your password to log in to the API or a desktop client. To receive or reset your security token, log in to Salesforce and click Setup > My Personal Information > Reset My Security Token . Password is case sensitive.
Service URL	URL of the Salesforce service you want to access. For example, <code>https://login.salesforce.com/services/Soap/u/50.0</code> In a test or development environment, you might want to access the Salesforce Sandbox testing environment. For more information about the Salesforce Sandbox, see the Salesforce documentation.
Refresh Token	Applicable for OAuth connection type. Refresh Token of Salesforce.
Consumer Key	Applicable for OAuth connection type. The Consumer Key obtained from Salesforce, required to generate the Refresh Token. For more information about how to generate the Consumer Key, see the Salesforce documentation.
Consumer Secret	Applicable for OAuth connection type. The Consumer Secret obtained from Salesforce, required to generate the Refresh Token. For more information about how to generate the Consumer Secret, see the Salesforce documentation.

Salesforce Marketing Cloud Connection Properties

Use a Salesforce Marketing Cloud connection to connect to a Salesforce Marketing Cloud object. You can create and manage a Salesforce Marketing Cloud connection in the Administrator tool or the Developer tool.

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

The following table describes the Salesforce Marketing Cloud connection properties:

Connection property	Description
Name	Name of the Salesforce Marketing Cloud connection.
ID	The Data Integration Service uses the ID to identify the connection.
Description	Optional. The description of the connection.
Location	Informatica domain where you want to create the connection.
Type	Connection type. Select Salesforce Marketing Cloud.

Connection property	Description
Salesforce Marketing Cloud Url	<p>The URL that the Data Integration Service uses to connect to the Salesforce Marketing Cloud WSDL.</p> <p>The following URL is an example for OAuth 1.0 URL: <code>https://webservice.s7.exacttarget.com/etframework.wsdl</code></p> <p>The following URL is an example for OAuth 2.0 URL: <code>https://<SUBDOMAIN>.soap.marketingcloudapis.com/etframework.wsdl</code></p> <p>Informatica recommends that you upgrade to OAuth 2.0 before Salesforce Marketing Cloud drops support for OAuth 1.0.</p>
Username	User name of the Salesforce Marketing Cloud account.
Password	Password for the Salesforce Marketing Cloud account.
ClientId	The client ID of Salesforce Marketing Cloud required to generate a valid access token.
ClientSecret	The client secret of Salesforce Marketing Cloud required to generate a valid access token.
Enable Logging	When you enable logging you can see the session log for the tasks.
UTC Offset	The Secure Agent uses the UTC offset connection property to read data from and write data to Salesforce Marketing Cloud in UTC offset time zone.
Batch Size	<p>Number of rows that the Secure Agent writes in a batch to the target.</p> <p>When you insert or update data and specify the contact key, the data associated with the specified contact ID is inserted or updated in a batch to Salesforce Marketing Cloud. When you upsert data to Salesforce Marketing Cloud, do not specify the contact key.</p>
Enable Multiple BU	Select this option if there are multiple business units in your Salesforce Marketing Cloud account. You can use the Salesforce Marketing Cloud connection to access data across all business units.

SAP Connection Properties

Use an SAP connection to access an SAP table or an SAP BW object. The SAP connection is an enterprise application connection. You can create and manage an SAP connection in the Administrator tool or the Developer tool.

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

The following table describes the SAP connection properties:

Property	Description
Username	Required. User name for the SAP source system that you want to access.
Password	Required. Password for the user name.
Connection type	<p>Required. Type of connection that you want to create.</p> <p>Select one of the following values:</p> <ul style="list-style-type: none"> - Application. Create an application connection when you want to connect to a specific SAP application server. - Load balancing. Create a load balancing connection when you want to use SAP load balancing. <p>Default is Application.</p> <p>Based on the connection type you select, the corresponding connection property fields become available in the Connection Details dialog box. The Developer tool greys out the connection property fields that are not applicable for a particular connection type.</p>
Host name	<p>Required when you create an SAP application connection.</p> <p>Host name or IP address of the SAP server that you want to connect to.</p>
System number	<p>Required when you create an SAP application connection.</p> <p>SAP system number.</p>
Message host name	<p>Required when you create an SAP load balancing connection.</p> <p>Host name of the SAP message server.</p>
R3 name/SysID	<p>Required when you create an SAP load balancing connection.</p> <p>Name of the SAP system.</p>
Group	<p>Required when you create an SAP load balancing connection.</p> <p>Group name of the SAP application server.</p>
Client	Required. SAP client number.
Language	<p>Optional. Language that you want to use for mappings and workflows.</p> <p>Must be compatible with the Developer tool code page.</p> <p>If you leave this option blank, the Developer tool uses the default language of the SAP system.</p>
Trace	<p>Optional. Use this option to track the JCo calls that the SAP system makes. SAP stores the information about the JCo calls in a trace file.</p> <p>Specify one of the following values:</p> <ul style="list-style-type: none"> - 0. Off - 1. Full <p>Default is 0.</p> <p>You can access the trace files from the following directories:</p> <ul style="list-style-type: none"> - <Informatica installation directory>/tomcat/bin directory on the machine where you installed the Informatica services - <Informatica installation directory>/clients/DeveloperClient directory on the machine where you installed the Developer tool
Additional parameters	<p>Optional. Enter any other connection parameter that you want to use.</p> <p>Use the following format:</p> <p><parameter name>=<value></p>

Property	Description
Staging directory	Path in the SAP system where the stage file will be created.
Source directory	Path that contains the source file. The path must be accessible to the Data Integration Service.
Use FTP	Enables FTP access to SAP.
FTP user	Required when you use FTP. User name to connect to the FTP server.
FTP password	Required when you use FTP. Password for the FTP user.
FTP host	Required when you use FTP. Host name or IP address of the FTP server. Optionally, you can specify a port number from 1 through 65535, inclusive. Default for FTP is 21. Use one of the following syntax to specify the host name: - hostname:port_number - IP address:port_number When you specify a port number, enable that port number for FTP on the host machine. If you enable SFTP, specify a host name or port number for an SFTP server. Default for SFTP is 22.
Retry period	Number of seconds that the Data Integration Service attempts to reconnect to the FTP host if the connection fails. If the Data Integration Service cannot reconnect to the FTP host in the retry period, the mapping or workflow fails. Default is 0. A value of 0 indicates an infinite retry period.
Use SFTP	Enables SFTP access to SAP.
Public key file name	Required when you enable SFTP and the SFTP server uses public key authentication. Public key file path and file name.
Private key file name	Required when you enable SFTP and the SFTP server uses public key authentication. Private key file path and file name.
Private key file name password	Required when you enable SFTP, and the SFTP server uses public key authentication and the private key is encrypted. Password to decrypt the private key file.
Port Range	HTTP port range that the Data Integration Service must use to read data from the SAP server in streaming mode. Enter the minimum and maximum port numbers with a hyphen as the separator. The minimum and maximum port number can range between 10000 and 65535. You can also specify the port range according to your organization. Default is 10000-65535.

Property	Description
Use HTTPS	Select this option to enable HTTPS streaming when you read data from SAP tables. By default, the Use HTTPS check box is not selected. For more information about configuring HTTPS for table reader mappings in streaming mode, see the article "HTTPS Configuration for Table Reader Mappings in Streaming Mode for PowerExchange for SAP NetWeaver" on the Informatica Documentation Portal.
Key store file path	Required when you use HTTPS. Path to the keystore file that contains the private or public key pairs and the associated certificates.
Key store password	Required when you use HTTPS. Password for the keystore file.
Private key password	Required when you use HTTPS. Password to decrypt the private key file.

Sequential Connection Properties

Use a Sequential connection to access sequential data sources. You create a Sequential connection in the Developer tool. You can manage a Sequential connection in the Administrator tool or the Developer tool.

A sequential data source is a data source that PowerExchange can access by using a data map defined with an access method of SEQ. The Data Integration Service connects to the data source through PowerExchange.

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

The following table describes Sequential connection properties:

Option	Description
Location	Node name for the location of the PowerExchange Listener that connects to the sequential data set. The node name is defined in the first parameter of the NODE statement in the PowerExchange dbmover.cfg configuration file.
User name	A user name that has the authority to access the sequential data set.

Option	Description
Password	<p>Password for the specified user name or a valid PowerExchange passphrase.</p> <p>A PowerExchange passphrase can be from 9 to 128 characters in length and can contain the following characters:</p> <ul style="list-style-type: none"> - Uppercase and lowercase letters - The numbers 0 to 9 - Spaces - The following special characters: ' - ; # \ , . / ! % & * () _ + { } : @ < > ? <p>Note: The first character is an apostrophe.</p> <p>Passphrases cannot include single quotation marks ('), double quotation marks ("), or currency symbols.</p> <p>To use passphrases, ensure that the PowerExchange Listener runs with a security setting of SECURITY=(1,N) or higher in the DBMOVER member. For more information, see "SECURITY Statement" in the <i>PowerExchange Reference Manual</i>.</p> <p>The allowable characters in the IBM IRRPHREX exit do not affect the allowable characters in PowerExchange passphrases.</p> <p>Note: A valid RACF passphrase can be up to 100 characters in length. PowerExchange truncates passphrases longer than 100 characters when passing them to RACF for validation.</p>
Code page	<p>Required. Name of the code page to use for reading from or writing to the sequential data set. Usually, this value is an ISO code page name, such as ISO-8859-6.</p>
Pass-through security enabled	<p>Enables pass-through security for the connection.</p>
Encryption type	<p>Optional. The type of encryption that the Data Integration Service uses. Select one of the following options:</p> <ul style="list-style-type: none"> - None - AES <p>Default is None.</p> <p>Note: Informatica recommends that you use Secure Sockets Layer (SSL) authentication instead of configuring the Encryption Type and Level connection properties. SSL authentication provides stricter security and is used by several Informatica products. For more information about implementing SSL authentication in a PowerExchange network, see the <i>PowerExchange Reference Manual</i>.</p>
[Encryption] Level	<p>If you select AES for Encryption Type, select one of the following options to indicate the encryption level that the Data Integration Service uses:</p> <ul style="list-style-type: none"> - 1. Use a 128-bit encryption key. - 2. Use a 192-bit encryption key. - 3. Use a 256-bit encryption key. <p>If you do not select AES for Encryption Type, this option is ignored.</p> <p>Default is 1.</p>
Pacing size	<p>Optional. Amount of data that the source system can pass to the PowerExchange Listener. Set the pacing size if an external application, database, or the Data Integration Service node is a bottleneck. User lower values for faster performance.</p> <p>Minimum value and default value is 0. A value of 0 provides the best performance.</p>
Interpret as rows	<p>Optional Select this option to express the pacing size as a number of rows. Clear this option to express the pacing size in kilobytes. By default, this option is not selected and the pacing size is in kilobytes.</p>

Option	Description
Compression	Optional. Select this option to enable source data compression. By compressing data, you can decrease the amount of data that Informatica applications send over the network. By default, this option is not selected and compression is disabled.
Offload processing	Optional. Controls whether to offload some bulk data processing from the source machine to the Data Integration Service machine. Select one of the following options: <ul style="list-style-type: none"> - AUTO. The Data Integration Service determines whether to use offload processing. - Yes. Use offload processing. - No. Do not use offload processing. Default is AUTO.
Worker threads	Optional. Number of threads that the Data Integration Service uses to process bulk data when offload processing is enabled. For optimal performance, this value should not exceed the number of available processors on the Data Integration Service machine. Valid values are 1 through 64. Default is 0, which disables multithreading.
Array size	Optional The number of records in the storage array for the worker threads. This option is applicable when you set the Worker Threads option to a value greater than 0. Valid values are 25 to 5000. Default is 25.
Write mode	Optional. Mode in which the Data Integration Service sends data to the PowerExchange Listener. Select one of the following write modes: <ul style="list-style-type: none"> - CONFIRMWRITEON. Sends data to the PowerExchange Listener and waits for a response before sending more data. Select this option when error recovery is a priority. However, this option might degrade performance. - CONFIRMWRITEOFF. Sends data to the PowerExchange Listener without waiting for a response. Use this option if you can reload the target table when an error occurs. - ASYNCHRONOUSWITHFAULTTOLERANCE. Sends data to the PowerExchange Listener without waiting for a response. This option also enables error detection. This option combines the speed of CONFIRMWRITEOFF and the data integrity of CONFIRMWRITEON. Default is CONFIRMWRITEON.

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.

Property	Description
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.
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>

Teradata Parallel Transporter Connection Properties

Use a Teradata PT connection to access Teradata tables. The Teradata PT connection is a database type connection. You can create and manage a Teradata PT connection in the Administrator tool or the Developer tool.

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

The following table describes Teradata PT 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.

Property	Description
Description	Description of the connection. The description cannot exceed 765 characters.
Location	Domain where you want to create the connection.
Type	Connection type. Select Teradata PT.
User Name	Teradata database user name with the appropriate read and write permissions to access the database.
Password	Password for the Teradata database user name.
Driver Name	Name of the Teradata JDBC driver.
Connection String	Connection string used to access metadata from the database. Use the following connection string: <code>jdbc:teradata://<hostname>/database=<database name>,tmode=ANSI,charset=UTF8</code>

The following table describes the properties for data access:

Property	Description
TDPID	Name or IP address of the Teradata database machine.
Database Name	Teradata database name. If you do not enter a database name, Teradata PT API uses the default login database name.
Data Code Page	Code page associated with the database. When you run a mapping that writes data to a Teradata target, the code page of the Teradata PT connection must be the same as the code page of the Teradata target. Default is UTF-8.
Tenacity	Number of hours that Teradata PT API continues trying to log on when the maximum number of operations run on the Teradata database. Must be a positive, non-zero integer. Default is 4.
Max Sessions	Maximum number of sessions that Teradata PT API establishes with the Teradata database. Must be a positive, non-zero integer. Default is 4.
Min Sessions	Minimum number of Teradata PT API sessions required for the Teradata PT API job to continue. Must be a positive integer between 1 and the Max Sessions value. Default is 1.
Sleep	Number of minutes that Teradata PT API pauses before it retries to log on when the maximum number of operations run on the Teradata database. Must be a positive, non-zero integer. Default is 6.
Use Metadata JDBC URL for TDCH	Indicates that the Teradata Connector for Hadoop (TDCH) must use the JDBC URL that you specified in the connection string under the metadata access properties. Default is selected. Clear this option to enter a different JDBC URL that TDCH must use when it runs the mapping.

Property	Description
TDCH JDBC Url	<p>Enter the JDBC URL that TDCH must use when it runs a Teradata mapping. Use the following format:</p> <pre>jdbc:teradata://<hostname>/database=<database name>,tmode=ANSI,charset=UTF8</pre> <p>This field is available only when you clear the Use Metadata JDBC URL for TDCH option.</p>
Data Encryption	<p>Enables full security encryption of SQL requests, responses, and data on Windows. Default is disabled.</p>
Additional Sqoop Arguments	<p>This property is applicable if you use a Hortonworks or Cloudera cluster, and run a Teradata mapping on the Blaze or Spark engine through Sqoop.</p> <p>Enter the arguments that Sqoop must use to process the data. For example, enter <code>--method split.by.amp</code>. Separate multiple arguments with a space.</p> <p>See the Hortonworks for Teradata Connector and Cloudera Connector Powered by Teradata documentation for a list of arguments that you can specify.</p> <p>Note: If you use Hortonworks for Teradata Connector, the <code>--split-by</code> argument is required if you add two or more source tables in the read operation. If you use Cloudera Connector Powered by Teradata, the <code>--split-by</code> argument is required in the source connection if the source table does not have a primary key defined.</p>
Authentication Type	<p>Method to authenticate the user.</p> <p>Select one of the following authentication types:</p> <ul style="list-style-type: none"> - Native. Authenticates your user name and password against the Teradata database specified in the connection. - LDAP. Authenticates user credentials against the external LDAP directory service. <p>Default is Native.</p>

Tableau Connection Properties

Use a Tableau connection to connect to Tableau. When you create a Tableau connection, you enter information to access Tableau.

The following table describes the Tableau connection properties:

Property	Description
Name	Name of the Tableau 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 Informatica domain where you want to create the connection.
Type	Type of connection. Select Tableau.

The following table describes the properties to connect to Tableau:

Connection Property	Description
Tableau Product	The name of the Tableau product to which you want to connect. You can choose one of the following Tableau products to publish the TDE or TWBX file: <ul style="list-style-type: none"> - Tableau Desktop. Creates a TDE file in the Data Integration Service machine. You can then manually import the TDE file to Tableau Desktop. - Tableau Server. Publishes the generated TDE or TWBX file to Tableau Server. - Tableau Online. Publishes the generated TDE or TWBX file to Tableau Online.
Connection URL	URL of Tableau Server or Tableau Online to which you want to publish the TDE or TWBX file. The URL has the following format: <code>http://<Host name of Tableau Server or Tableau Online>:<port></code>
User Name	User name of the Tableau Server or Tableau Online account.
Password	Password for the Tableau Server or Tableau Online account.
Content URL	The name of the site on Tableau Server or Tableau Online where you want to publish the TDE or TWBX file. Contact the Tableau administrator to provide the site name.

Tableau V3 Connection Properties

When you set up a Tableau V3 connection, you must configure the connection properties.

The following table describes the Tableau V3 connection properties:

Property	Description
Name	Name of the Tableau V3 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 Informatica domain where you want to create the connection.
Type	Type of connection. Select Tableau V3.

The following table describes the properties to connect to Tableau:

Connection Property	Description
Tableau Product	<p>The name of the Tableau product to which you want to connect.</p> <p>You can choose one of the following Tableau products to publish the <code>.hyper</code> or TWBX file:</p> <p>Tableau Desktop</p> <p>Creates a <code>.hyper</code> file in the Data Integration Service machine. You can then manually import the <code>.hyper</code> file to Tableau Desktop.</p> <p>Tableau Server</p> <p>Publishes the generated <code>.hyper</code> or TWBX file to Tableau Server.</p> <p>Tableau Online</p> <p>Publishes the generated <code>.hyper</code> or TWBX file to Tableau Online.</p>
Connection URL	<p>The URL of Tableau Server or Tableau Online to which you want to publish the <code>.hyper</code> or TWBX file.</p> <p>Enter the URL in the following format: <code>http://<Host name of Tableau Server or Tableau Online>:<port></code></p>
User Name	The user name of the Tableau Server or Tableau Online account.
Password	The password for the Tableau Server or Tableau Online account.
Site ID	<p>The ID of the site on Tableau Server or Tableau Online where you want to publish the or TWBX file.</p> <p>Note: Contact the Tableau administrator to provide the site ID.</p>
Schema File Path	<p>The path to a sample <code>.hyper</code> file from where the Data Integration Service imports the Tableau metadata.</p> <p>Enter one of the following options for the schema file path:</p> <ul style="list-style-type: none"> - Absolute path to the <code>.hyper</code> file. - Directory path for the <code>.hyper</code> files. - Empty directory path. <p>The path you specify for the schema file becomes the default path for the target <code>.hyper</code> file. If you do not specify a file path, the Data Integration Service uses the following default file path for the target <code>.hyper</code> file:</p> <pre><Data Integration Service installation directory>/apps/ Data_Integration_Server/<latest version>/bin/rtdm</pre>

Twitter Streaming Connection Properties

Use a Twitter Streaming connection to access near-real time data from the Twitter web site. The Twitter Streaming connection is a connection to the social media company's streaming API. You can create and manage a Twitter Streaming connection in the Administrator tool or the Developer tool.

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

The following table describes the general properties for a Twitter Streaming connection:

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.
Type	The connection type. Select Twitter Streaming.

The following table describes the properties for hose type and OAuth authentication:

Property	Description
Hose Type	Streaming API methods. You can specify one of the following methods: - Filter. The Twitter <code>statuses/filter</code> method returns public statuses that match the search criteria. - Sample. The Twitter <code>statuses/sample</code> method returns a random sample of all public statuses.
Consumer Key	The consumer key that you get when you create the application in Twitter. Twitter uses the key to identify the application.
Consumer Secret	The consumer secret that you get when you create the Twitter application. Twitter uses the secret to establish ownership of the consumer key.
Do you have OAuth details?	Indicates whether you want to configure OAuth. Select one of the following values: - Yes. Indicates that you have the access token and secret. - No. Launches the OAuth Utility.
Access Token	Access token that the OAuth Utility returns. Twitter uses the token instead of the user credentials to access the protected resources.
Access Secret	Access secret that the OAuth Utility returns. The secret establishes ownership of a token.

VSAM Connection Properties

Use a VSAM connection to access VSAM data tables. The VSAM connection is a flat file connection type. You create a VSAM connection in the Developer tool. You can manage a VSAM connection in the Administrator tool or the Developer tool.

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

The following table describes VSAM connection properties:

Option	Description
Location	Node name for the location of the PowerExchange Listener that connects to the VSAM data set. The node name is defined in the first parameter of the NODE statement in the PowerExchange dbmover.cfg configuration file.
User name	A user name that has the authority to connect to the VSAM data set.
Password	<p>A password for the specified user or a valid PowerExchange passphrase.</p> <p>A PowerExchange passphrase can be from 9 to 128 characters in length and can contain the following characters:</p> <ul style="list-style-type: none"> - Uppercase and lowercase letters - The numbers 0 to 9 - Spaces - The following special characters: ' - ; # \ , . / ! % & * () _ + { } : @ < > ? <p>Note: The first character is an apostrophe.</p> <p>Passphrases cannot include single quotation marks ('), double quotation marks ("), or currency symbols.</p> <p>To use passphrases, ensure that the PowerExchange Listener runs with a security setting of SECURITY=(1,N) or higher in the DBMOVER member. For more information, see "SECURITY Statement" in the <i>PowerExchange Reference Manual</i>.</p> <p>The allowable characters in the IBM IRRPHREX exit do not affect the allowable characters in PowerExchange passphrases.</p> <p>Note: A valid RACF passphrase can be up to 100 characters in length. PowerExchange truncates passphrases longer than 100 characters when passing them to RACF for validation.</p>
Code page	Required. Name of the code page to use for reading from or writing to the VSAM data set. Usually, this value is an ISO code page name, such as ISO-8859-6.
Pass-through security enabled	Enables pass-through security for the connection.
Encryption type	<p>Optional. The type of encryption that the Data Integration Service uses. Select one of the following options:</p> <ul style="list-style-type: none"> - None - AES <p>Default is None.</p> <p>Note: Informatica recommends that you use Secure Sockets Layer (SSL) authentication instead of configuring the Encryption Type and Level connection properties. SSL authentication provides stricter security and is used by several Informatica products. For more information about implementing SSL authentication in a PowerExchange network, see the <i>PowerExchange Reference Manual</i>.</p>
[Encryption] Level	<p>If you select AES for Encryption Type, select one of the following options to indicate the encryption level that the Data Integration Service uses:</p> <ul style="list-style-type: none"> - 1. Use a 128-bit encryption key. - 2. Use a 192-bit encryption key. - 3. Use a 256-bit encryption key. <p>If you do not select AES for Encryption Type, this option is ignored.</p> <p>Default is 1.</p>

Option	Description
Pacing size	Optional. Amount of data that the source system can pass to the PowerExchange Listener. Set the pacing size if an external application, database, or the Data Integration Service node is a bottleneck. User lower values for faster performance. Minimum value and default value is 0. A value of 0 provides the best performance.
Interpret as rows	Optional. Select this option to express the pacing size as a number of rows. Clear this option to express the pacing size in kilobytes. By default, this option is not selected and the pacing size is in kilobytes.
Compression	Optional. Select this option to enable source data compression. By compressing data, you can decrease the amount of data that Informatica applications send over the network. By default, this option is not selected and compression is disabled.
Offload processing	Optional. Controls whether to offload some bulk data processing from the source machine to the Data Integration Service machine. Select one of the following options: <ul style="list-style-type: none"> - AUTO. The Data Integration Service determines whether to use offload processing. - Yes. Use offload processing. - No. Do not use offload processing. Default is AUTO.
Worker threads	Optional. Number of threads that the Data Integration Service uses to process bulk data when offload processing is enabled. For optimal performance, this value should not exceed the number of available processors on the Data Integration Service machine. Valid values are 1 through 64. Default is 0, which disables multithreading.
Array size	Optional. The number of records in the storage array for the worker threads. This option is applicable when you set the Worker Threads option to a value greater than 0. Valid values are 25 to 5000. Default is 25.
Write mode	Optional. Mode in which Data Integration Service sends data to the PowerExchange Listener. Select one of the following write modes: <ul style="list-style-type: none"> - CONFIRMWRITEON. Sends data to the PowerExchange Listener and waits for a response before sending more data. Select this option when error recovery is a priority. However, this option might degrade performance. - CONFIRMWRITEOFF. Sends data to the PowerExchange Listener without waiting for a response. Use this option if you can reload the target table when an error occurs. - ASYNCHRONOUSWITHFAULTTOLERANCE. Sends data to the PowerExchange Listener without waiting for a response. This option also enables error detection. This option combines the speed of CONFIRMWRITEOFF and the data integrity of CONFIRMWRITEON. Default is CONFIRMWRITEON.

Web Services Connection Properties

Use a web services connection to connect a Web Service Consumer transformation to a web service.

The following table describes the web services connection properties:

Property	Description
Username	User name to connect to the web service. Enter a user name if you enable HTTP authentication or WS-Security. If the Web Service Consumer transformation includes WS-Security ports, the transformation receives a dynamic user name through an input port. The Data Integration Service overrides the user name defined in the connection.
Password	Password for the user name. Enter a password if you enable HTTP authentication or WS-Security. If the Web Service Consumer transformation includes WS-Security ports, the transformation receives a dynamic password through an input port. The Data Integration Service overrides the password defined in the connection.
End Point URL	URL for the web service that you want to access. The Data Integration Service overrides the URL defined in the WSDL file. If the Web Service Consumer transformation includes an endpoint URL port, the transformation dynamically receives the URL through an input port. The Data Integration Service overrides the URL defined in the connection.
Timeout	Number of seconds that the Data Integration Service waits for a response from the web service provider before it closes the connection. Specify a timeout value between 1 and 10,000 seconds.
HTTP Authentication Type	Type of user authentication over HTTP. Select one of the following values: <ul style="list-style-type: none">- None. No authentication.- Automatic. The Data Integration Service chooses the authentication type of the web service provider.- Basic. Requires you to provide a user name and password for the domain of the web service provider. The Data Integration Service sends the user name and the password to the web service provider for authentication.- Digest. Requires you to provide a user name and password for the domain of the web service provider. The Data Integration Service generates an encrypted message digest from the user name and password and sends it to the web service provider. The provider generates a temporary value for the user name and password and stores it in the Active Directory on the Domain Controller. It compares the value with the message digest. If they match, the web service provider authenticates you.- NTLM. Requires you to provide a domain name, server name, or default user name and password. The web service provider authenticates you based on the domain you are connected to. It gets the user name and password from the Windows Domain Controller and compares it with the user name and password that you provide. If they match, the web service provider authenticates you. NTLM authentication does not store encrypted passwords in the Active Directory on the Domain Controller.

Property	Description
WS Security Type	Type of WS-Security that you want to use. Select one of the following values: <ul style="list-style-type: none"> - None. The Data Integration Service does not add a web service security header to the generated SOAP request. - PasswordText. The Data Integration Service adds a web service security header to the generated SOAP request. The password is stored in the clear text format. - PasswordDigest. The Data Integration Service adds a web service security header to the generated SOAP request. The password is stored in a digest form which provides effective protection against replay attacks over the network. The Data Integration Service combines the password with a nonce and a time stamp. The Data Integration Service applies a SHA hash on the password, encodes it in base64 encoding, and uses the encoded password in the SOAP header.
Trust Certificates File	File containing the bundle of trusted certificates that the Data Integration Service uses when authenticating the SSL certificate of the web service. Enter the file name and full directory path. Default is <Informatica installation directory>/services/shared/bin/ca-bundle.crt.
Client Certificate File Name	Client certificate that a web service uses when authenticating a client. Specify the client certificate file if the web service needs to authenticate the Data Integration Service.
Client Certificate Password	Password for the client certificate. Specify the client certificate password if the web service needs to authenticate the Data Integration Service.
Client Certificate Type	Format of the client certificate file. Select one of the following values: <ul style="list-style-type: none"> - PEM. Files with the .pem extension. - DER. Files with the .cer or .der extension. Specify the client certificate type if the web service needs to authenticate the Data Integration Service.
Private Key File Name	Private key file for the client certificate. Specify the private key file if the web service needs to authenticate the Data Integration Service.
Private Key Password	Password for the private key of the client certificate. Specify the private key password if the web service needs to authenticate the Data Integration Service.
Private Key Type	Type of the private key. PEM is the supported type.

Identifier Properties in Database Connections

When you create most relational database connections, you must configure database identifier properties. The identifier properties determine whether the Data Integration Service encloses identifiers within delimited characters when the service generates SQL queries to access the database.

A database identifier is a database object name. Tables, views, columns, indexes, triggers, procedures, constraints, and rules can have identifiers. You use the identifier to reference the object in SQL queries. A database can have regular identifiers or delimited identifiers that must be enclosed within delimited characters.

Regular Identifiers

Regular identifiers comply with the format rules for identifiers. Regular identifiers do not require delimited characters when they are used in SQL queries.

For example, the following SQL statement uses the regular identifiers *MYTABLE* and *MYCOLUMN*:

```
SELECT * FROM MYTABLE
WHERE MYCOLUMN = 10
```

Delimited Identifiers

Delimited identifiers must be enclosed within delimited characters because they do not comply with the format rules for identifiers.

Databases can use the following types of delimited identifiers:

Identifiers that use reserved keywords

If an identifier uses a reserved keyword, you must enclose the identifier within delimited characters in an SQL query. For example, the following SQL statement accesses a table named *ORDER*:

```
SELECT * FROM "ORDER"
WHERE MYCOLUMN = 10
```

Identifiers that use special characters

If an identifier uses special characters, you must enclose the identifier within delimited characters in an SQL query. For example, the following SQL statement accesses a table named *MYTABLE\$@*:

```
SELECT * FROM "MYTABLE$@"
WHERE MYCOLUMN = 10
```

Case-sensitive identifiers

By default, identifiers in IBM DB2, Microsoft SQL Server, and Oracle databases are not case sensitive. Database object names are stored in uppercase, but SQL queries can use any case to refer to them. For example, the following SQL statements access the table named *MYTABLE*:

```
SELECT * FROM mytable
SELECT * FROM MyTable
SELECT * FROM MYTABLE
```

To use case-sensitive identifiers, you must enclose the identifier within delimited characters in an SQL query. For example, the following SQL statement accesses a table named *MyTable*:

```
SELECT * FROM "MyTable"
WHERE MYCOLUMN = 10
```

Identifier Properties

When you create most database connections, you must configure database identifier properties. The identifier properties that you configure depend on whether the database uses regular identifiers, uses keywords or special characters in identifiers, or uses case-sensitive identifiers.

Configure the following identifier properties in a database connection:

SQL Identifier Character

Type of character that the database uses to enclose delimited identifiers in SQL queries. The available characters depend on the database type.

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.

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.

Support Mixed-case Identifiers

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.

In the Informatica client tools, you must refer to the identifiers with the correct case. For example, when you create the database connection, you must enter the database user name with the correct case.

When the **SQL Identifier Character** property is set to none, the **Support Mixed-case Identifiers** property is disabled.

Example: Database Uses Regular Identifiers

In this example, the database uses regular identifiers. No identifiers contain a reserved keyword or a special character. The database uses identifiers that are not case sensitive.

In the database connection, set the **SQL Identifier Character** property to (None). When **SQL Identifier Character** is set to none, the **Support Mixed-case Identifiers** property is disabled.

When the Data Integration Service generates SQL queries, the service does not place delimited characters around any identifiers.

Example: Database Uses Keywords or Special Characters in Identifiers

In this example, the database uses keywords or special characters in some identifiers. The database uses identifiers that are not case sensitive.

In the database connection, configure the identifier properties as follows:

1. Set the **SQL Identifier Character** property to the character that the database uses for delimited identifiers.

This example sets the property to `"` (*quotes*).

2. Clear the **Support Mixed-case Identifiers** property.

When the Data Integration Service generates SQL queries, the service places the selected character around identifiers that use a reserved keyword or that use a special character. For example, the Data Integration Service generates the following query:

```
SELECT * FROM "MYTABLES@" /* identifier with special characters enclosed within
delimited
character */
WHERE MYCOLUMN = 10 /* regular identifier not enclosed within delimited character */
```

Example: Database Uses Case-Sensitive Identifiers

In this example, the database uses case-sensitive identifiers. The database might use keywords or special characters in some identifiers, or it might not.

In the database connection, configure the identifier properties as follows:

1. Set the **SQL Identifier Character** property to the character that the database uses for delimited identifiers.

This example sets the property to `"` (*quotes*).

2. Select the **Support Mixed-case Identifiers** property.

When the Data Integration Service generates SQL queries, the service places the selected character around all identifiers. For example, the Data Integration Service generates the following query:

```
SELECT * FROM "MyTable"      /* case-sensitive identifier enclosed within delimited
character */
WHERE "MYCOLUMN" = 10      /* regular identifier enclosed within delimited character */
```

CHAPTER 10

Schedules

This chapter includes the following topics:

- [Schedules Overview, 213](#)
- [Create and Edit Schedules, 213](#)
- [Pausing and Resuming a Schedule, 216](#)
- [Removing Jobs from a Schedule, 217](#)
- [Deleting a Schedule, 217](#)

Schedules Overview

You can create a schedule to indicate when to run deployed mappings and workflows. You can also schedule the date and time to run profiles and scorecards.

When you create a schedule, you configure when the schedule runs. You can add jobs to the schedule when you create it or you can save the schedule and add jobs to it later.

You might define a schedule to automate routine tasks or manage resource usage in the domain. For example, you can schedule large mappings to run at different times to prevent the mappings from overloading the nodes. You can also pause the schedule to prevent the jobs from running during holidays or during domain maintenance.

Create and Edit Schedules

Before you schedule a job, create a schedule that defines when the job runs. Jobs can run one time or at an interval. You can schedule jobs when you create the schedule or you can add jobs to the schedule later.

You can select a time zone when you configure the date and time to run the schedule. By default, when you enter a start time for a schedule, the time zone is the time zone of your client machine. If the Data Integration Service runs on a machine in a different time zone from the client, you need to ensure that the job runs at a specific time for that time zone. You can choose the time zone when you define the schedule.

You can schedule deployed mappings and workflows including mappings and workflows for profiles and scorecards. Optionally, you can configure a parameter file or parameter set for the job. Configure a parameter file or set when you want to run the deployed mapping or workflow with different settings. If the Scheduler Service runs on multiple nodes, you must store the parameter file in a directory that all of the nodes can access. Configure the directory in the Storage Properties for the Scheduler Service.

Optionally, you can configure the **Run As User** property for the job. You might configure the property if you need to schedule an object that you do not have permission on. When you configure the **Run As User** property, you run the job as a user in the domain who has permission on the object.

After you create a schedule, you can edit the schedule. When you edit a schedule, you can change the recurrence options, schedule jobs, or change the job parameters.

Creating a Schedule

Create a schedule to run deployed mappings, deployed workflows, profiles, or scorecards at a specified time.

1. In the Administrator tool, click the **Manage** tab.
2. Click the **Schedules** view.
3. Click **Actions > Create Schedule**.
The **Properties** dialog box appears.
4. Enter a name and an optional description for the schedule.
5. In the **Recurrence** section, choose Run Once to run the schedule one time, or choose Recurring Schedule to run the schedule on an interval.
6. If you are creating a recurring schedule, specify whether to run the job on a Daily, Weekly, or Monthly interval.

The following table describes the recurrence options that you can configure:

Option	Description
Daily - Run every	Run the job every n days
Daily - Run after every	Run the job every n minutes or hours
Weekly	Run the job on a day or days of the week
Monthly - Run every n day	Run the job on a day of the month
Monthly - Run every	Run the job on a day even n week of the month
Monthly - Run every last day of the month	Run the job on the last day of the month

7. In the **Start** section, configure a start date and time for the schedule.
8. Select the time zone for the schedule from the list.
The default time zone is the time zone of the client machine.
9. If you are creating a recurring schedule, you can optionally configure an end date and time.
You can end the recurrence on a specific date or after n runs.

The following image shows the Scheduler Wizard:

Properties - Step 1 of 3

General

Name *

Description

Recurrence

Run once

Recurring Schedule

Daily

Weekly

Monthly

Run every days

Run after every

Start

Time Zone

Start Date

Start Time HH:MM

End

No End Date

End Date

End Time HH:MM

End After : Runs

10. Click **Next**.
The **Applications** dialog box appears.
11. Expand a Data Integration Service and select jobs to run on the schedule.
12. Optionally, upload a parameter file to define parameters for the mapping or workflow.
 - a. Select a mapping or a workflow.
The **Application Properties** section appears.
 - b. Click **Upload**.
The **Upload Parameter File** dialog box appears.
 - c. Click **Choose Files**.
 - d. Browse for a file, and then click **Open**.
 - e. Click **OK**.
13. Optionally, select a parameter set.
14. Optionally, configure the **Run As User** property to run the job as a different user.
 - a. Click **Change**.
The **Run As** dialog box appears.
 - b. Expand the user list.

- c. Select a user.
 - d. Enter the password for the user.
 - e. Click **OK**.
15. Click **Next**.
- The **Review** window appears and lists the schedule properties.
16. Click **Finish** to create the schedule.
- The schedule appears in the list of schedules in the Navigator.

Editing a Schedule

You can edit a schedule in the Administrator tool. You can edit a schedule to change recurrence options, schedule jobs, or run jobs with different parameters.

1. In the Administrator tool, click the **Manage** tab.
2. Click the **Schedules** view.
The Domain Navigator lists the schedules in the domain.
3. In the Domain Navigator, select a schedule.
The contents panel shows the schedule properties.
4. To edit the schedule, click **Actions > Edit Schedule**.

Pausing and Resuming a Schedule

You can pause and resume a schedule from the Administrator tool. For example, you might want to pause a schedule when you perform maintenance on the Scheduler Service.

When you pause a schedule, the jobs that run on the schedule stop running until you resume the schedule.

1. In the Administrator tool, click the **Manage** tab.
2. Click the **Schedules** view.
The **All Schedules** view lists the schedules in the domain.
3. Select a schedule. Use the **Ctrl** and **Shift** keys to select multiple schedules.
4. To pause the schedule, click **Pause**.
The **Pause Confirmation** message appears.
5. Click **OK**.
The **Schedule Status Change** message appears.
6. Click **OK**.
The schedule status changes to Paused.
7. To resume the schedule, select the schedule, and then click **Resume**.
The **Resume Confirmation** message appears.
8. Click **OK**.
The **Schedule Status Change** message appears.

9. Click **OK**.
The schedule status changes to Scheduled.

Removing Jobs from a Schedule

You can remove a deployed mapping, profile, scorecard, or workflow job from a schedule.

1. In the Administrator tool, click the **Manage** tab.
2. Click the **Schedules** view.
3. Click the **Scheduled Jobs** view.
A list of scheduled jobs appears in the contents panel.
4. Select the job that you would like to remove from a schedule.
5. Click **Actions > Remove Schedule Association**.
The **Remove Job Confirmation** message displays.
6. Click **OK**.
The job is removed from the schedule and the **Scheduled Jobs** view.

Deleting a Schedule

You can delete a schedule in the Administrator tool.

When you delete a schedule, any jobs that are running finish running, and all future runs are canceled. If you do not want the jobs to stop running, add them to a different schedule before you delete the schedule.

1. In the Administrator tool, click the **Manage** tab.
2. Click the **Schedules** view.
The **All Schedules** view lists all of the schedules in the domain.
3. In the **All Schedules** view, select a schedule. Use the **Ctrl** and **Shift** keys to select multiple schedules.
4. Click **Actions > Delete Schedule**.
The **Delete Confirmation** dialog box appears.
5. Click **OK** to delete the schedule.

CHAPTER 11

Domain Object Export and Import

This chapter includes the following topics:

- [Domain Object Export and Import Overview, 218](#)
- [Export Process, 218](#)
- [View Domain Objects, 219](#)
- [Import Process, 220](#)

Domain Object Export and Import Overview

You can use the command line to migrate objects between two different domains of the same version.

You might migrate domain objects from a development environment to a test or production environment.

To export and import domain objects, use the following infacmd isp commands:

ExportDomainObjects

Exports native users, native groups, roles, and connections to an XML file.

ImportDomainObjects

Imports native users, native groups, roles, and connections into an Informatica domain.

You can use an infacmd control file to filter the objects during the export or import.

You can also use the infacmd xrf generateReadableViewXML command to generate a readable XML file from an export file. You can review the readable XML file to determine if you need to filter the objects that you import.

Export Process

You can use the command line to export domain objects from a domain.

Perform the following tasks to export domain objects:

1. Determine the domain objects that you want to export.
2. If you do not want to export all domain objects, create an export control file to filter the objects that are exported.
3. Run the infacmd isp exportDomainObjects command to export the domain objects.

The command exports the domain objects to an export file. You can use this file to import the objects into another domain.

Rules and Guidelines for Exporting Domain Objects

Review the following rules and guidelines before you export domain objects:

- When you export a user, by default, you do not export the user password. If you do not export the password, the administrator must reset the password for the user after the user is imported into the domain. However, when you run the `infacmd isp exportDomainObjects` command, you can choose to export an encrypted version of the password.
- When you export a user, you do not export the associated groups of the user. If applicable, assign the user to the group after you import the user and group.
- When you export a group, you export all sub-groups and users in the group.
- You cannot export the Administrator user, the Administrator role, the Everyone group, or LDAP users or groups. To replicate LDAP users and groups in an Informatica domain, import the LDAP users and groups directly from the LDAP directory service.
- To export native users and groups from domains of different versions, use the `infacmd isp exportUsersAndGroups` command.
- When you export a connection, by default, you do not export the connection password. If you do not export the password, the administrator must reset the password for the connection after the connection is imported into the domain. However, when you run the `infacmd isp exportDomainObjects` command, you can choose to export an encrypted version of the password.

View Domain Objects

You can view domain object names and properties in the export XML file.

Run `infacmd xrf generateReadableViewXML` command, to create a readable XML from the export file.

The following section provides a sample readable XML file:

```
<global:View xmlns:global="http://global" xmlns:connection="http://connection"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" xsi:schemaLocation="
  http://connection connection.xsd http://global globalSchemaDomain.xsd http://global
  globalSchema.xsd">
  <NativeUser isAdmin="false" name="admin" securityDomain="Native" viewId="0">
    <UserInfo email="" fullName="admin" phone="" viewId="1"/>
  </NativeUser>
  <User isAdmin="false" name="User1" securityDomain="Native" viewId="15">
    <UserInfo email="" fullName="NewUSer" phone="" viewId="16"/>
  </User>
  <Group name="TestGroup1" securityDomain="Native" viewId="182">
    <UserRef name="User1" securityDomain="Native" viewId="183"/>
    <UserRef name="User6" securityDomain="Native" viewId="188"/>
  </Group>
  <Role customRole="false" name="Administrator" viewId="242">
    <Description viewId="243">Provides all privilege and permission access to an
  Informatica service.</Description>
    <ServicePrivilegeDefinition name="PwxListenerService" viewId="244">
      <Privilege category="" isEnabled="true" name="close" viewId="245"/>
      <Privilege category="" isEnabled="true" name="closeforce" viewId="246"/>
      <Privilege category="" isEnabled="false" name="Management Commands" viewId="249"/>
      <Privilege category="" isEnabled="false" name="Informational Commands"
  viewId="250"/>
    </ServicePrivilegeDefinition>
  </Role>
```

```
<Connection connectionString="inqa85sql25@qa90"
connectionType="SQLServerNativeConnection"
domainName="" environmentsSQL="" name="conn4" ownerName=""
schemaName="" transactionSQL="" userName="dummy" viewId="7512">
<ConnectionPool maxIdleTime="120" minConnections="0" usePool="true" viewId="7514"/>
</Connection>
</global:View>
```

Viewable Domain Object Names

You can view the following domain object names in the readable XML file:

User
UserInfo
Role
ServicePrivilegeDef
Privilege
Group
GroupRef
UserRef
ConnectInfo
ConnectionPoolAttributes

Supported Connection Types

DB2iNativeConnection
DB2NativeConnection
DB2zNativeConnection
JDBCConnection
ODBCNativeConnection
OracleNativeConnection
PWXMetaConnection
SAPConnection

Import Process

You can use the command line to import domain objects from an export file into a domain.

Perform the following tasks to import domain objects:

1. Run the `infacmd xrf generateReadableViewXML` command to generate a readable XML file from an export file. Review the domain objects in the readable XML file and determine the objects that you want to import.
2. If you do not want to import all domain objects in the export file, create an import control file to filter the objects that are imported.
3. Run the `infacmd isp importDomainObjects` command to import the domain objects into the specified domain.

4. After you import the objects, you may still have to create other domain objects such as application services and folders.

Rules and Guidelines for Importing Domain Objects

Review the following rules and guidelines before you import domain objects:

- When you import a group, you import all sub-groups and users in the group.
- To import native users and groups from domains of different versions, use the `infacmd isp importUsersAndGroups` command.
- After you import a user or group, you cannot rename the user or group.
- You import roles independently of users and groups. Assign roles to users and groups after you import the roles, users, and groups.
- You cannot import the Administrator group, the Administrator user, the Administrator role, the Everyone group, or LDAP users or groups.

Conflict Resolution

A conflict occurs when you try to import an object with a name that exists for an object in the target domain. Configure the conflict resolution to determine how to handle conflicts during the import.

You can define a conflict resolution strategy through the command line or control file when you import the objects. The control file takes precedence if you define conflict resolution in the command line and control file. The import fails if there is a conflict and you did not define a conflict resolution strategy.

You can configure one of the following conflict resolution strategies:

Reuse

Reuses the object in the target domain.

Rename

Renames the source object. You can provide a name in the control file, or else the name is generated. A generated name has a number appended to the end of the name.

Replace

Replaces the target object with the source object.

Merge

Merges the source and target objects into one group. For example, if you merge groups with the same name, users and sub-groups from both groups are merged into the group in the target domain.

You cannot define the merge conflict resolution strategy through the command line. Use a control file to define the merge conflict resolution strategy. You must include the group object type section with merge as the conflict resolution policy with reuse, replace, or rename for all conflicting users in the control file.

For example, specify the merge conflict resolution strategy for the following groups:

- Group A with users a1, a2, b1, b2 in the source domain.
- Group A with users a1, a2, a3, b1, b2 in the target domain

You get the following results in the group after merge in the target domain:

- a1, a2, b1, b2 if you choose reuse or replace
- a1, a2, a3, b1, b2 if you choose rename.

CHAPTER 12

License Management

This chapter includes the following topics:

- [License Management Overview, 222](#)
- [Types of License Keys, 224](#)
- [Creating a License Object, 225](#)
- [Assigning a License to a Service, 226](#)
- [Unassigning a License from a Service, 226](#)
- [Updating a License, 227](#)
- [Removing a License, 227](#)
- [License Properties, 228](#)

License Management Overview

The Service Manager on the master gateway node manages Informatica licenses.

A license enables you to perform the following tasks:

- Run application services, such as the Analyst Service, Data Integration Service, and PowerCenter Repository Service.
- Use add-on options, such as partitioning for PowerCenter, grid, and high availability.
- Access particular types of connections, such as Oracle, Teradata, Microsoft SQL Server, and IBM MQ Series.
- Use Metadata Exchange options, such as Metadata Exchange for Cognos and Metadata Exchange for Rational Rose.

When you install Informatica, the installation program creates a license object in the domain based on the license key that you used during installation.

You assign a license object to each application service to enable the service. For example, you must assign a license to the PowerCenter Integration Service before you can use the PowerCenter Integration Service to run a workflow.

You can create additional license objects in the domain. Based on your project requirements, you may need multiple license objects. For example, you may have two license objects, where each license object allows you to run services on a different operating system. You might also use multiple license objects to manage multiple projects in the same domain. One project may require access to particular database types, while the other project does not.

License Validation

The Service Manager validates application service processes when they start. The Service Manager validates the following information for each service process:

- Product version. Verifies that you are running the appropriate version of the Informatica services.
- Platform. Verifies that the Informatica services are running on a licensed operating system.
- Expiration date. Verifies that the license is not expired. If the license expires, no application service assigned to the license can start. You must assign a valid license to the Informatica services to start them.
- PowerCenter options. Determines the options that the Informatica services have permission to use. For example, the Service Manager verifies if the PowerCenter Integration Service can use the Session on Grid option.
- Connectivity. Verifies connections that the Informatica services have permission to use. For example, the Service Manager verifies that PowerCenter can connect to a IBM DB2 database.
- Metadata Exchange options. Determines the Metadata Exchange options that are available for use. For example, the Service Manager verifies that you have access to the Metadata Exchange for Business Objects Designer.

Licensing Log Events

The Service Manager generates log events and writes them to the Log Manager. It generates log events for the following actions:

- You create or delete a license.
- You apply an incremental license key to a license.
- You assign an application service to a license.
- You unassign a license from an application service.
- The license expires.
- The Service Manager encounters an error, such as a validation error.

The log events include the user name and the time associated with the event.

You must have permission on the domain to view the logs for Licensing events.

The Licensing events appear in the domain logs.

License Management Tasks

You can perform the following tasks to manage the licenses:

- Create the license in the Administrator tool. You use a license key to create a license in the Administrator tool.
- Assign a license to each application service. Assign a license to each application service to enable the service.
- Unassign a license from an application service. Unassign a license from an application service if you want to discontinue the service or migrate the service from a development environment to a production environment. After you unassign a license from a service, you cannot enable the service until you assign another valid license to it.
- Update the license. Update the license to add PowerCenter options to the existing license.
- Remove the license. Remove a license if it is obsolete.

- Configure user permissions on a license.
- View license details. You may need to review the licenses to determine details, such as expiration date and the maximum number of licensed CPUs. You may want to review these details to ensure you are in compliance with the license. Use the Administrator tool to determine the details for each license.
- Monitor license usage and licensed options. You can monitor the usage of logical CPUs and PowerCenter Repository Services. You can monitor the number of software options purchased for a license and the number of times a license exceeds usage limits in the License Management Report.

You can perform all of these tasks in the Administrator tool or by using *infacmd isp* commands.

Types of License Keys

Informatica provides license keys in license files. The license key is encrypted. When you create the license from the license key file, the Service Manager decrypts the license key and enables the purchased options.

You create a license from a license key file. You apply license keys to the license to enable additional options. Informatica uses the following types of license keys:

- Original keys. Informatica generates an original key based on your contract. Informatica may provide multiple original keys depending on your contract.
- Incremental keys. Informatica generates incremental keys based on updates to an existing license, such as an extended license period or an additional option.

Note: Informatica licenses typically change with each version. Use a license key file valid for the current version to ensure that your installation includes all functionality.

Original Keys

Original keys identify the contract, product, and licensed features. Licensed features include the Informatica edition, deployment type, number of authorized CPUs, and authorized Informatica options and connectivity. You use the original keys to install Informatica and create licenses for services. You must have a license key to install Informatica. The installation program creates a license object for the domain in the Administrator tool. You can use other original keys to create more licenses in the same domain. You use a different original license key for each license object.

Incremental Keys

You use incremental license keys to update an existing license. You add an incremental key to an existing license to add or remove options, such as PowerCenter options, connectivity, and Metadata Exchange options. For example, if an existing license does not allow high availability, you can add an incremental key with the high availability option to the existing license.

The Service Manager updates the license expiration date if the expiration date of an incremental key is later than the expiration date of an original key. The Service Manager uses the latest expiration date. A license object can have different expiration dates for options in the license. For example, the IBM DB2 relational connectivity option may expire on 12/01/2006, and the session on grid option may expire on 04/01/06.

The Service Manager validates the incremental key against the original key used to create the license. An error appears if the keys are not compatible.

Creating a License Object

You can create a license object in a domain and assign the license to application services. You can create the license in the Administrator tool using a license key file. The license key file contains an encrypted original key. You use the original key to create the license.

You can also use the *infacmd isp* AddLicense command to add a license to the domain.

Use the following guidelines to create a license:

- Use a valid license key file. The license key file must contain an original license key. The license key file must not be expired.
- You cannot use the same license key file for multiple licenses. Each license must have a unique original key.
- Enter a unique name for each license. You create a name for the license when you create the license. The name must be unique among all objects in the domain.
- Put the license key file in a location that is accessible by the Administrator tool computer. When you create the license object, you must specify the location of the license key file.

After you create the license, you can change the description. To change the description of a license, select the license in Navigator of the Administrator tool, and then click Edit.

1. In the Administrator tool, click **Actions > New > License**.
The **Create License** window appears.
2. Enter the following options:

Option	Description
Name	Name of the license. The name is not case sensitive and must be unique within the domain. It cannot exceed 128 characters or begin with @. It also cannot contain spaces or the following special characters: ` ~ % ^ * + = { } \ ; : ' " / ? . , < > ! () []
Description	Description of the license. The description cannot exceed 765 characters.
Path	Path of the domain in which you create the license. Read-only field. Optionally, click Browse and select a domain in the Select Folder window. Optionally, click Create Folder to create a folder for the domain.
License File	File containing the original key. Click Browse to locate the file.

If you try to create a license using an incremental key, a message appears that states you cannot apply an incremental key before you add an original key.

You must use an original key to create a license.

3. Click **Create**.

Assigning a License to a Service

Assign a license to an application service before you can enable the service. When you assign a license to a service, the Service Manager updates the license metadata. You can also use the *infacmd isp AssignLicense* command to assign a license to a service.

1. Select the license in the **Domain Navigator** of the Administrator tool.
2. Click the **Assigned Services** tab.
3. In the **License** tab, click **Actions > Edit Assigned Services**.

The **Assign or Unassign this license to the services** window appears.

4. Select the services under **Unassigned Services**, and click **Add**.
Use Ctrl-click to select multiple services. Use Shift-click to select a range of services. Optionally, click **Add all** to assign all services.
5. Click **OK**.

Rules and Guidelines for Assigning a License to a Service

Use the following rules and guidelines when you assign licenses:

- You can assign licenses to disabled services.
- If you want to assign a license to a service that has a license assigned to it, you must first unassign the existing license from the service.
- To start a service with backup nodes, you must assign it to a license with high availability.
- To restart a service automatically, you must assign the service to a license with high availability.

Unassigning a License from a Service

You might need to unassign a license from a service if the service becomes obsolete or if you want to discontinue a service. You might want to discontinue a service if you are using more CPUs than you are licensed to use.

You can use the Administrator tool or the *infacmd isp UnassignLicense* command to unassign a license from a service.

You must disable a service before you can unassign a license from it. After you unassign the license from the service, you cannot enable the service. You must assign a valid license to the service to reenable it.

You must disable the service before you can unassign the license. If you try to unassign a license from an enabled service, a message appears that states you cannot remove the service because it is running.

1. Select the license in the **Domain Navigator** of the Administrator tool.
2. Click the **Assigned Services** tab.
3. In the **License** tab, click **Actions > Edit Assigned Services**.

The **Assign or Unassign this license to the services** window appears.

4. Select the service under **Assigned Services**, and then click **Remove**. Optionally, click **Remove all** to unassign all assigned services.
5. Click **OK**.

Updating a License

You can update the current license in the Informatica domain with an incremental license key..

When you add an incremental key to a license, the Service Manager adds or removes licensed options and updates the license expiration date.

You can also use the `infacmd isp UpdateLicense` command to add an incremental key to a license.

After you update the license you must restart Informatica services for the changes to take effect.

Use the following guidelines to update a license:

- Verify that the license key file is accessible by the Administrator tool computer. When you update the license object, you must specify the location of the license key file.
- The incremental key must be compatible with the original key. An error appears if the keys are not compatible.

The Service Manager validates the license key against the original key based on the following information:

- Serial number
- Deployment type
- Distributor
- Informatica edition
- Informatica version

1. In the Administrator tool, click the **Manage** tab > **Services and Nodes** view.
2. Select a **license** in the Domain Navigator.
3. Click **Manage Actions** > **Add Incremental Key**.
The **Update License** window appears.
4. Click **Choose Files** to browse for a license key file.
5. Click **OK**.
6. In the **License Details** section of the **Properties** view, click **Edit** to edit the description of the license.
7. Click **OK**.

Removing a License

You can remove a license from a domain using the Administrator tool or the `infacmd isp RemoveLicense` command.

Before you remove a license, disable all services assigned to the license. If you do not disable the services, all running service processes abort when you remove the license. When you remove a license, the Service Manager unassigns the license from each assigned service and removes the license from the domain. To re-enable a service, assign another license to it.

If you remove a license, you can still view License Usage logs in the Log Viewer for this license, but you cannot run the License Report on this license.

To remove a license from the domain:

1. Select the license in the **Domain Navigator** of the Administrator tool.

2. Click **Actions > Delete**.

License Properties

You can view license details using the Administrator tool or the `infacmd isp ShowLicense` command.

The license details are based on all license keys applied to the license. The Service Manager updates the existing license details when you add a new incremental key to the license.

You might review license details to determine options that are available for use. You may also review the license details and license usage logs when monitoring licenses.

For example, you can determine the number of CPUs your company is licensed to use for each operating system.

To view license details, select the license in the **Domain Navigator**.

The Administrator tool displays the license properties in the following sections:

- **License Details.** View license details on the **Properties** tab. Shows license attributes, such as the license object name, description, and expiration date.
- **Supported Platforms.** View supported platforms on the **Properties** tab. Shows the operating systems and how many CPUs are supported for each operating system.
- **Repositories.** View the licensed repositories on the **Properties** tab. Shows the maximum number of licensed repositories.
- **Assigned Services.** View application services that are assigned to the license on the **Assigned Services** tab.
- **PowerCenter Options.** View the PowerCenter options on the **Options** tab. Shows all licensed PowerCenter options, such as session on grid, high availability, and pushdown optimization.
- **Connections.** View the licensed connections on the **Options** tab. Shows all licensed connections. The license enables you to use connections, such as DB2 and Oracle database connections.
- **Metadata Exchange Options.** View the Metadata Exchange options on the **Options** tab. Shows a list of all licensed Metadata Exchange options, such as Metadata Exchange for Business Objects Designer.

You can also run the License Management Report to monitor licenses.

License Details

You can use the license details to view high-level information about the license. Use this license information when you audit the licensing usage.

The general properties for the license appear in the **License Details** section of the **Properties** tab.

The following table describes the general properties for a license:

Property	Description
Name	Name of the license.
Description	Description of the license.

Property	Description
Location	Path to the license in the Navigator.
Edition	PowerCenter Advanced edition.
License Version	Version of license.
Distributed By	Distributor of the product.
Issued On	Date when the license was issued to the customer.
Expires On	Date when the license expires.
Validity Period	Period for which the license is valid.
Serial Number	Serial number of the license. The serial number identifies the customer or project. If you have multiple PowerCenter installations, there is a separate serial number for each project. The original and incremental keys for a license have the same serial number.
Deployment Level	Level of deployment. Values are "Development" and "Production."

You can also use the license event logs to view audit summary reports. You must have permission on the domain to view the logs for license events.

Supported Platforms

You assign a license to each service. The service can run on any operating system supported by the license. One product license can support multiple operating system platforms.

The supported platforms for the license appear in the Supported Platforms section of the **Properties** tab.

The following table describes the supported platform properties for a license:

Property	Description
Description	Name of the supported operating system.
Logical CPUs	Number of CPUs you can run on the operating system.
Issued On	Date on which the license was issued.
Expires	Date on which the license expires.

Repositories

The maximum number of active repositories for the license appear in the Repositories section of the Properties tab.

The following table describes the repository properties for a license:

Property	Description
Description	Name of the repository.
Instances	Number of repository instances running on the operating system.
Issued On	Date on which the license was issued for this option.
Expires	Date on which the license expires for this option.

Service Options

The license enables you to use Informatica Service options such as data cleansing, data federation, and pushdown optimization.

The options for the license appear in the Service Options section of the **Options** tab.

Connections

The license enables you to use connections such as DB2 and Oracle database connections. The license also enables you to use connections for PowerExchange adapters such as PowerExchange for Facebook.

The connections for the license appear in the Connections section of the **Options** tab.

Metadata Exchange Options

The license enables you to use Metadata Exchange options such as Metadata Exchange for Business Objects Designer and Metadata Exchange for Microstrategy.

The Metadata Exchange options for the license appear in the Metadata Exchange Options section of the **Options** tab.

CHAPTER 13

Monitoring

This chapter includes the following topics:

- [Monitoring Overview, 231](#)
- [Configuring Monitoring, 232](#)
- [Optimizing Monitoring Performance, 234](#)
- [Summary Statistics, 235](#)
- [Monitor Data Integration Services , 237](#)
- [Monitor Ad Hoc Jobs, 238](#)
- [Monitor Applications, 242](#)
- [Monitor Deployed Mapping Jobs, 243](#)
- [Monitor Logical Data Objects, 246](#)
- [Monitor SQL Data Services, 247](#)
- [Monitor Web Services, 250](#)
- [Monitor Workflows, 251](#)
- [Job Status After Application Service Restart or Failover, 259](#)
- [Monitoring a Folder of Objects, 259](#)

Monitoring Overview

In Informatica Administrator, you can monitor statistics for Data Integration Service jobs on the **Monitor** tab.

After you configure the monitoring Model repository at the domain level, you can view monitoring statistics in the Administrator tool. The **Monitor** tab displays current and historical information about multiple Data Integration Service and integration objects. Use the **Summary Statistics** view to view graphical summaries of object state and distribution across multiple Data Integration Service. You can also view graphs of the memory and CPU that the Data Integration Services used to run the objects. Use the **Execution Statistics** view to monitor properties, run-time statistics, and run-time reports for integration objects.

You can monitor the following objects:

- Ad hoc jobs
- Applications
- Logical data objects
- SQL data services

- Web services
- Workflows

Monitoring is a domain function that the Service Manager performs. The Service Manager stores the monitoring configuration in the monitoring Model repository. The Service Manager also persists, updates, retrieves, and publishes run-time statistics for integration objects in the monitoring Model repository.

You can also access monitoring from the following tools:

The Informatica Monitoring tool

The Monitoring tool is a direct link to the **Monitor** tab of the Administrator tool. The Monitoring tool is useful if you do not need access to other Administrator tool features. You must have at least one monitoring privilege to access the Monitoring tool. You can access the Monitoring tool from the following URL:

```
http://<Administrator tool host><Administrator tool port>/monitoring/
```

The Analyst tool

You can monitor objects on the **Job Status** tab in the Analyst tool. The **Job Status** tab shows the status of Analyst tool jobs, such as profile jobs, scorecard jobs, and jobs that load mapping specification results to the target.

The Developer tool

You can open the Monitoring tool from the Developer tool. When you monitor from the Developer tool, you can view jobs that users run from the Developer tool. The Monitoring tool shows the status of Developer tool jobs, such as mapping jobs.

Configuring Monitoring

You can configure the monitoring Model Repository Service at the domain level. After you configure the Monitoring Configuration parameters, the Monitor tab displays the statistics and reports about the objects in the domain. Statistics and reports appear in the **History** view on the **Manage** tab and in the **Summary Statistics** and **Execution Statistics** views on the **Monitor** tab.

The monitoring Model Repository Service stores statistics and reports for the Data Integration Service jobs. The statistics include historical information about objects that the multiple Data Integration Services run. The reports show key metrics about integration objects.

If you do not configure monitoring, then some views on the **Manage** and **Monitor** tabs do not have content. The workflow graph is also empty, and notifications disappear when you refresh the page.

To view monitoring statistics and reports, perform the following tasks:

1. Configure monitoring settings. Configure a Model repository as a monitoring Model repository to store run-time statistics for objects that the multiple Data Integration Services runs.
2. Configure reports and statistics views. Choose which statistics appear in the **Statistics** and **Reports** views.

Note: In a domain that uses Kerberos authentication, users must have the Administrator role for the monitoring Model Repository Service that stores statistics. If users do not have Administrator role, some statistics might not appear.

Step 1. Configure Monitoring Settings

Configure monitoring settings for the domain when you want to view historical information about the domain. When you configure monitoring settings, you specify a Model repository as the monitoring Model repository to store run-time statistics for Data Integration Service jobs.

Create the monitoring Model repository content before you configure the monitoring settings. If you create the content after you configure the monitoring settings, you must recycle the monitoring Model Repository Service after the content is created.

1. In the Administrator tool, click the **Manage** tab > **Services and Nodes** view.
2. In the Domain Navigator, select the domain.
3. In the Domain section, click the **Monitoring Configuration** view.
The current monitoring configuration appears.
4. Click **Edit** to change the monitoring configuration.
5. Edit the following options:

Option	Description
Model Repository Service	Name of the Model repository that stores the historical information. The Model repository must not be integrated with a version control system.
Username	User name to access the Model Repository Service. Does not appear in domains that use Kerberos authentication.
Password	Password of the user name to access the Model Repository Service. Does not appear in domains that use Kerberos authentication.
Modify Password	Modify the Model Repository Service password.
Security Domain	Name of the security domain that the Model repository user belongs to.
Preserve Summary Historical Data	Number of days that the Model repository saves averaged data. If purging is disabled, then the Model repository saves the data indefinitely. Default is 180. Minimum is 0. Maximum is 366.
Preserve Detailed Historical Data	Number of days that the Model repository saves per-minute data. If purging is disabled, then the Model repository saves the data indefinitely. Default is 14. Minimum is 1. Maximum is 14.
Purge Statistics Every	Interval, in days, at which the Model Repository Service purges data that is older than the values configured in the Preserve Historical Data option. Default is 1 day.
Days At	Time of day when the Model Repository Service purges statistics. Default is 1:00 a.m.
Maximum Number of Sortable Records	Maximum number of records that can be sorted in the Monitor tab. If the number of records on the Monitor tab is greater than this value, then you can only sort by Start Time and End Time . Default is 3,000.
Maximum Delay for Update Notifications	Maximum time, in seconds, that the Data Integration Service buffers statistics before it stores them in the Model repository and displays them in the Monitor tab. If the Data Integration Service shuts down unexpectedly before it stores the statistics in the Model repository, then the statistics are lost. Default is 10.

Option	Description
Show Milliseconds in Date Time Field	Include milliseconds for date and time fields in the Monitor tab.

6. Click **OK**.

To apply the settings, you must restart all of the Data Integration Services.

Step 2. Configure Reports and Statistics Views

By default, the **Statistics** and **Reports** views on the **Execution Statistics** view are empty. To view statistics and reports, you must configure the Report and Statistic settings in the domain. These settings apply to all Data Integration Services in the domain.

Before you configure statistics and reports, you must specify a monitoring Model Repository Service in the Monitoring Configuration tab and enable the monitoring Model Repository Service.

1. In the Administrator tool, click the **Monitor > Execution Statistics** tab.
2. Click **Actions > Report and Statistic Settings**.
3. In the **Report and Statistic Settings** dialog box, click the **Statistics** tab.
4. Configure the time ranges that you want to use for statistics, and select the frequency at which the statistics assigned to each time range should be updated.
5. In the **Default Time Range** list, select a default time range to appear for all statistics.
6. Click the **Reports** tab.
7. Enable the time ranges that you want to use for reports, and select the frequency at which the reports assigned to each time range should be updated.
8. In the **Default Time Range** list, select a default time range to appear for all reports.
9. Click **Select Reports**.
10. In the **Select Reports** dialog box, add the reports that you want to run to the **Selected Reports** box.
11. Organize the reports in the order in which you want to view them on the **Monitor** tab.
12. Click **OK** to close the **Select Reports** dialog box.
13. Click **OK** in the **Report and Statistic Settings** dialog box.

Optimizing Monitoring Performance

When you configure monitoring in the domain, you configure the monitoring Model repository to store monitoring statistics. You also configure how often the monitoring Model Repository Service purges statistics. You can configure the repository database and purge settings to minimize resource consumption and maximize monitoring performance.

To optimize monitoring performance, enable the **Purge Statistics** option in the Monitoring Configuration. When you enable statistics purging, you configure how often the monitoring Model Repository Service purges data. Use the **Preserve** options in the Monitoring Configuration to configure how long Detailed and Summary data are saved in the monitoring Model repository before they are purged.

For optimum monitoring performance, consider the following guidelines when you configure the domain for monitoring:

- Create the monitoring Model repository to store monitoring data. Specify the monitoring Model Repository Service when you configure monitoring in the domain.
- Configure the monitoring Model Repository Service on the machine where you configure the domain.
- Verify that the monitoring Model repository database is tuned as required for the following properties:
 - Maximum Heap Size property
 - Java Stack Size property
 - Memory settings
 - Hibernate Connection Pool Size property
- Purge monitoring data during off-hours to limit the impact on other database operations.
- Purge monitoring data daily.

Summary Statistics

The **Summary Statistics** view shows information about the Data Integration Services and the objects that a Data Integration Service runs.

Use the **Summary Statistics** view to view summaries of object distribution, object state, and Data Integration Service resource usage for a specified time range. You can view statistics for the domain, or by Data Integration Service or application.

You can view statistics about the following objects:

- Ad hoc jobs. Jobs that users ran from the Developer tool or the Analyst tool.
- Deployed mappings and workflows. Mappings and workflows that are deployed in an application.
- Requests and connections. Deployed SQL data services and web services.
- Resource Usage. CPU and memory usage of the Data Integration Service processes that are running in the domain or on a node in the domain.

For example, you can view all of the jobs in the domain that failed over the last eight hours. You can review resource usage for that time range to see if the jobs failed due to resource issues on the node.

The following image shows the **Summary Statistics** view with a list of ad hoc jobs in the **Details** panel:

Name	Type	Service Name	State	ID	Started By	Start Time	Elapsed Time	End Time
CUSTOMER	Preview	DIS1	Completed	Kz_QzD...	Administrator	06/29/2016 17:04:09	00:00:08	06/29/2016 17:04:18
VT_NATION	Preview	DIS1	Completed	eCl8pD...	Administrator	06/29/2016 17:34:57	00:00:05	06/29/2016 17:35:02
Resource:GetMapping	Preview	DIS1	Completed	Tm3NE...	Administrator	06/30/2016 08:50:02	00:00:18	06/30/2016 08:50:20

When you select a source and time range you can choose one of the following options to view the data:

- Graphical distribution. Displays doughnut and line charts that summarize object distribution and resource usage for a time range. The doughnut charts show distributions of objects by state, type, and the Data Integration Service that ran them. The line charts compare resource usage for the Data Integration Service that ran the jobs to resource usage for all of the processes that ran on the node during that time range.
- Tabular distribution. View the total completed, running, canceled, aborted, and failed jobs.
- Details. View a list of the jobs, requests, or connections that comprise the summary statistics. You can click the Job ID to view the specific job in the Execution Statistics.
- Export data. Export the detail data for a specific object type to a .csv file.

The **Summary Statistics** view displays statistics using data that is stored in the monitoring Model repository. You must configure a monitoring Model Repository Service in the Monitoring Configuration tab before you can view **Summary Statistics**.

Viewing Summary Statistics

You can view summary and detail information about ad hoc jobs, deployed mappings and workflows, requests and connections, and resource usage for Data Integration Services in the domain.

1. Click the **Monitor** tab.
2. Click the **Summary Statistics** view.
The time line appears.
3. In the **Source** field, choose the source for which you want to view statistics.
4. In the **Time Range** field, choose the time range for which you want to view statistics.
5. Optionally, choose **Custom** to specify a custom time range.
6. Click **View Statistics**.
The object and resource usage panels appear.
7. Expand an object type to view statistics for that object type.
8. Select whether to view a graphical distribution, a tabular distribution, or a detail list of the data. Or, choose to export a .csv file.

9. Optionally, you can perform the following actions in the **Resource Usage** panel:

Option	Description
Show	Show all nodes or one node.
Click the magnifying glass	Enlarge the chart.
Drag	Zoom in on a section of the chart.
Reset Zoom	View the chart at the default size.

Monitor Data Integration Services

You can monitor Data Integration Services on the **Execution Statistics** view on the **Monitor** tab.

When you select a Data Integration Service in the Navigator, the contents panel shows the following views:

- **Properties** view
- **Reports** view

Properties View for a Data Integration Service

The **Properties** view shows the general properties and run-time statistics for objects that the selected Data Integration Service ran.

When you select a Data Integration Service in the Navigator, you can view the general properties and run-time statistics.

General Properties for a Data Integration Service

You can view general properties, such as the service name, object type, and description. The **Persist Statistics Enabled** property indicates whether the Data Integration Service stores persisted statistics in the monitoring Model repository. This option is true when you configure the global settings for the domain.

You can also view information about objects that the Data Integration Service runs. To view information about an object, select the object in the Navigator or contents panel. Depending on the object type, details about the object appear in the contents panel or details panel.

Statistics for a Data Integration Service

You can view run-time statistics about objects that the Data Integration Service runs. Select the object type and time period to display the statistics. You can view statistics about jobs, applications, connections, requests, and workflows. For example, you can view the number of failed, canceled, and completed profiling jobs in the last four hours.

Reports View for a Data Integration Service

The **Reports** view shows reports about objects that the selected Data Integration Service runs.

When you select a Data Integration Service in Navigator of the **Monitor** tab, the **Reports** view shows reports about jobs. For example, you can view the **Most Active Users for Jobs** report to determine users that ran the most jobs during a specific time period. Click a link in the report to show more details about the objects

included in the link. For example, you can click the number of failed deployed mappings to see details about each deployed mapping that failed.

Monitor Ad Hoc Jobs

You can monitor ad hoc jobs on the **Monitor** tab. Ad hoc jobs are jobs that users run from the Developer tool or the Analyst tool.

An ad hoc job is a preview, scorecard, profile, mapping, audit, or reference table process that a user runs from the Developer tool or the Analyst tool. When a user runs a job, a Data Integration Service runs the job process and the job appears in the Monitor tab.

You can run up to five jobs at a time from the Developer tool. All remaining jobs are queued, and do not appear in the Monitor tab until they run.

By default, you can monitor jobs that you run. If you have the appropriate monitoring privilege, you can also view jobs that other users run.

When you select **Ad Hoc Jobs** in the Navigator of the **Execution Statistics** view, a list of jobs appears in the contents panel. The contents panel groups related jobs based on the job type. You can expand a job type to view the related jobs under it. For example, when you run a profile job, the Data Integration Service converts the job into a mapping. The mapping appears under the profile job in the contents panel.

When you select a job in the contents panel, you can view logs for the job, view context for the job, or cancel the job. You can also view properties for the job in the details panel. Depending on the type of job, the details panel might show general properties, mapping properties, or statistics.

When you select Ad Hoc jobs in the Navigator of the Execution Statistics view, a list of jobs appears in the contents panel. The contents panel groups related jobs based on the job type. You can expand a job type to view the related jobs under it.

The following list describes the types of properties and statistics that can appear in the details panel:

Properties

Shows the general properties about the selected job, such as the name, job type, user who ran the job, and start time of the job. If the job ran on a grid, the details panel shows the node that the job ran on.

Mapping Properties

You can view mapping properties when you select a profile or scorecard job in the contents panel. These jobs have an associated mapping. You can view mapping properties such as the request ID, the mapping name, and the log file name. You can also view throughput and resource usage statistics for the associated mappings.

Blaze Execution Plan

You can view a Blaze execution plan when you run a mapping with the Blaze engine in the Hadoop environment. The Blaze execution plan displays the Blaze engine script that the Data Integration Service generates based on the mapping logic, the unique identifier for the script, and the tasks that the script depends on.

Summary Statistics

You can view summary statistics you select an ad hoc mapping job, deployed mapping job, or mapping object in a workflow in the contents panel. The **Summary Statistics** view displays throughput and resource usage statistics for the job run.

The following image shows the **Summary Statistics** view for a mapping job:

MappingLookup						
		Properties		Summary Statistics		Detailed Statistics
▼ Throughput						
Source	Rows	Average Rows/Sec	Bytes	Average Bytes/Sec	First Row Accessed	Dropped Rows
Read_CUSTOMER_DE...	4001	4001	392098	392098	09/04/2015 12:30:17	0
Target						
Write_CUSTOMER_DETAILS...	Rows	Average Rows/Sec	Bytes	Average Bytes/Sec	Rejected Rows	
Write_Customer_Details...	4001	4001	424106	424106	0	
Write_Flat_File_Data_Object	4001	4001	16004	16004	0	
▼ Resource Usage						
Executing Node			node_715			
Average CPU Usage			0 %			
Average Memory Usage			53 MB			

You can also view the summary statistics for the source or target on the **Summary Statistics** tab, such as processed rows and bytes.

Note: Only the **number of rows processed** appears in the Summary Statistics for the Hive source and target. The remaining property values appears as 0 or N/A for the Hive source or target.

Detailed Statistics

You can view detailed statistics for an ad hoc mapping job, deployed mapping job, or mapping object in a workflow in the contents panel. Detailed statistics appear for jobs that run longer than one minute in separate local processes. The **Detailed Statistics** view displays graphs of the throughput and resource usage statistics for the job run.

The following image shows the **Detailed Statistics** view for a mapping job in a workflow:

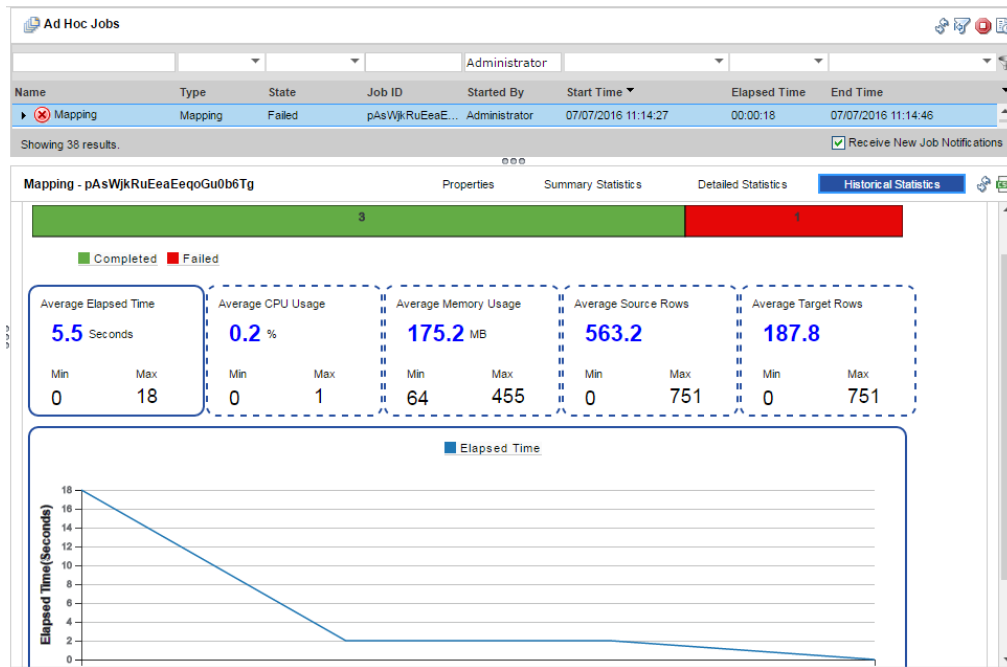


Historical Statistics

You can view historical statistics when you select an ad hoc mapping job, deployed mapping job, or mapping object in a workflow in the **Contents** panel. The **Historical Statistics** view shows averaged data from the last 500 runs for a specific job. For example, you can view the minimum, maximum, and average duration of the mapping job. You can view the average amount of CPU that the job consumes

when it runs. You can choose to view averages for mapping jobs with different states, such as completed, canceled, or aborted. A bar graph shows the number of jobs in each state. Click a link below the bar graph to view statistics for a specific state.

The following image shows the **Historical Statistics** view for a mapping job that completed three times and failed one time:



Aggregated Cluster Logs

You can get aggregated cluster logs for deployed Hadoop mappings, Databricks mappings, auto-deploy cluster jobs, local files, and Spark data preview jobs based on the job ID. You can get a .zip or tar.gz file of the aggregated cluster logs for a job and write the compressed aggregated log file to a target directory. You can also use the `infacmd ms fetchAggregatedClusterLogs` command or use the REST API to collect the aggregated cluster logs for a job.

Name	Type	State	Job ID	Started By	Start Time	Elapsed ...	End Time
Mapping_f...	Mapping	Completed	D_KDU3...	Administr...	05/22/2019 15:06:23	00:01:29	05/22/2019 15:07:52
MAINS...	Grid Task	Completed	D_KDU3...	Administr...	05/22/2019 15:06:23	00:01:27	05/22/2019 15:07:51
Mapping_f...	Mapping	Completed	Bv1cHxv...	Administr...	05/22/2019 14:23:11	00:00:31	05/22/2019 14:23:42
Mapping_f...	Mapping	Completed	wu9zQ3x...	Administr...	05/22/2019 14:21:17	00:00:35	05/22/2019 14:21:52
Mapping_f...	Mapping	Completed	T-wXq3x...	Administr...	05/22/2019 14:18:04	00:00:48	05/22/2019 14:18:52
Mapping_f...	Mapping	Canceled	BGrd-3xu...	Administr...	05/22/2019 14:15:57	00:00:01	05/22/2019 14:15:58
Mapping_f...	Mapping	Completed	s7j88HxV...	Administr...	05/22/2019 11:21:54	00:03:18	05/22/2019 11:25:12

For more information about the `infacmd ms fetchAggregatedClusterLogs` command, see the *Informatica 10.4.0 Command Reference*.

Viewing Logs for an Ad Hoc Job

You can download the logs for a job to view the job details.

1. In the Administrator tool, click the **Monitor** tab.
2. Click the **Execution Statistics** view.
3. In the Domain Navigator, expand a Data Integration Service and select **Ad Hoc Jobs**.
4. In the contents panel, select a job.
5. Click **Actions > View Logs for Selected Object**.

A dialog box appears with the option to open or save the log file.

Canceling an Ad Hoc Job

You can cancel a running job. You may want to cancel a job that hangs or that is taking an excessive amount of time to complete.

1. In the Administrator tool, click the **Monitor** tab.
2. Click the **Execution Statistics** view.
3. In the Domain Navigator, expand a Data Integration Service and select **Ad Hoc Jobs**.
4. In the contents panel, select a job.
5. Click **Actions > Cancel Selected Object**.

Viewing Summary Statistics for an Ad Hoc Job

You can view throughput and resource usage statistics for ad hoc mapping jobs.

1. In the Administrator tool, click the **Monitor** tab.
 2. Click the **Execution Statistics** view.
 3. In the Domain Navigator, expand a Data Integration Service and select **Ad Hoc Jobs**.
- A list of jobs appears in the contents panel.
4. In the contents panel, select a job.
- The details panel displays the Properties for the job.
5. Click the **Summary Statistics** view in the details panel.

The **Summary Statistics** view displays throughput and resource usage statistics for the source and target.

Optionally, you can sort the statistics in ascending or descending order. Click a column header to sort the column in ascending order. Click the column header again to sort the column in descending order.

Viewing Detailed Statistics for an Ad Hoc Job

You can view graphs of the throughput and resource usage for ad hoc mapping jobs that run in separate local processes. Detailed statistics appear for jobs that run longer than one minute.

1. In the Administrator tool, click the **Monitor** tab.
 2. Click the **Execution Statistics** view.
 3. In the Domain Navigator, expand a Data Integration Service and select **Ad Hoc Jobs**.
- A list of jobs appears in the contents panel.

4. In the contents panel, select a job.
The details panel displays the Properties for the job.
5. Click the **Detailed Statistics** view in the details panel.
The **Detailed Statistics** view displays the throughput graph and resource usage graphs.

Optionally, you can complete the following tasks in the **Detailed Statistics** view:

Task	Description
Enlarge a graph	Move the cursor over a graph, and then click the magnifying glass icon.
Enlarge a section of an enlarged graph	Drag the cursor to select an area to enlarge.
Switch between rows and bytes in the throughput graph	Click the Bytes option or the Rows option.
Choose which statistics are plotted on the throughput graph	In the throughput field, select the sources and targets that you want to view.

Monitoring Mapping Audits

When you audit a data engineering mapping, you can monitor the audit job as an ad hoc job.

Audit jobs are listed with the type *Audit Mapping*.

If an audit configuration contains both rules that run before the mapping and rules that run after the mapping, separate audit jobs run for the pre-mapping and post-mapping jobs.

Monitor Applications

You can monitor applications on the **Monitor** tab.

When you select an application in the Navigator of the **Execution Statistics** view, the contents panel shows the following views:

- **Properties** view
- **Reports** view

You can expand an application in the Navigator to monitor the application components.

Properties View for an Application

The **Properties** view shows general properties and run-time statistics about each application and the objects in an application. Applications can include deployed mapping jobs, logical data objects, SQL data services, web services, and workflows.

When you select an application in the Navigator of the **Execution Statistics** view, you can view general properties and run-time statistics.

General Properties

You can view general properties, such as the name and description of the application. You can also view additional information about the objects in an application. To view information about an object, select the folder in the Navigator and the object in the contents panel. Details about the object appear in the details panel.

Statistics

You can view run-time statistics about an application and about the jobs, connections, requests, and workflows associated with the application. For example, you can view the number of enabled and disabled applications, number of aborted connections, and number of completed, failed, and canceled jobs and workflows.

Reports View for an Application

The **Reports** view shows monitoring reports about the selected application.

When you select an application in the Navigator of the **Execution Statistics** view, the **Reports** view shows reports about the application components.

Monitor Deployed Mapping Jobs

You can monitor deployed mapping jobs on the **Execution Statistics** view on the **Monitor** tab, or in the Monitoring tool.

You can view information about deployed mapping jobs in an application.

You can monitor a deployed mapping in the following locations:

- Monitoring tool. In the Developer tool, click the **Menu** button in the Progress view and select **Monitor Jobs**. Select the Data Integration Service that runs the mapping and click **OK**. The Monitoring tool opens. Expand an application in the Navigator and select the **Deployed Mapping Jobs** folder. A list of deployed mapping jobs appears in the contents panel. You can view mapping execution statistics in the Monitoring tool. The REST Operations Hub generates statistics based on the API configuration.
- Administrator tool. Expand an application in the Navigator and select the **Deployed Mapping Jobs** folder. A list of deployed mapping jobs appears in the contents panel.

The contents panel shows properties about each deployed mapping job, such as Job ID, name of the mapping, state of the job, and start time of the job. If you run the job on a grid, the contents panel also shows the node that the Data Integration Service that runs the process is running on.

Select a deployed mapping job in the contents panel to view logs for the job, reissue the job, cancel the job, or view statistics about the job. You can view statistics about throughput and resource usage for the job run.

Viewing Logs for a Deployed Mapping Job

You can download the logs for a deployed mapping job to view the job details.

Note: The log contents for a deployed mapping job depend on how the Data Integration Service is configured. For more information about logs when a Data Integration Service grid is configured to run jobs in separate remote processes, see the *Informatica Application Service Guide*.

1. In the Administrator tool, click the **Monitor** tab.

2. Click the **Execution Statistics** view.
3. In the Domain Navigator, expand a Data Integration Service.
4. Expand an application and select **Deployed Mapping Jobs**.
A list of mapping jobs appear in the contents panel.
5. Select a mapping job.
6. Click **Actions > View Logs for Selected Object**.
A dialog box appears with the option to open or save the log file.

Reissuing a Deployed Mapping Job

You can reissue a deployed mapping job when the mapping jobs fails. When you reissue a deployed mapping job, the Data Integration Service runs the job again.

1. Click the **Execution Statistics** view.
2. In the Domain Navigator, expand a Data Integration Service.
3. Expand an application and select **Deployed Mapping Jobs**.
The contents panel displays a list of deployed mapping jobs.
4. Select a deployed mapping job.
5. Click **Actions > Reissue Selected Object**.

Canceling a Deployed Mapping Job

You can cancel a deployed mapping job. You might want to cancel a deployed mapping job that hangs or that is taking an excessive amount of time to complete.

1. In the Administrator tool, click the **Monitor** tab.
2. Click the **Execution Statistics** view.
3. In the Domain Navigator, expand a Data Integration Service.
4. Expand an application and select **Deployed Mapping Jobs**.
The contents panel displays a list of deployed mapping jobs.
5. Select a deployed mapping job.
6. Click **Actions > Cancel Selected Job**.

Viewing Summary Statistics for a Deployed Mapping Job

You can view throughput and resource usage statistics for deployed mapping jobs.

1. In the Administrator tool, click the **Monitor** tab.
2. Click the **Execution Statistics** view.
3. In the Domain Navigator, expand a Data Integration Service.
4. Expand an application and select **Deployed Mapping Jobs**.
A list of mapping jobs appears in the contents panel.
5. Select a mapping job.
The details panel displays the Properties for the mapping job.

6. Click the **Summary Statistics** view.

The **Summary Statistics** view displays throughput and resource usage statistics for the source and target.

Optionally, you can sort the statistics in ascending or descending order. Click a column header to sort the column in ascending order. Click the column header again to sort the column in descending order.

Viewing Detailed Statistics for a Deployed Mapping Job

You can view graphs of the throughput and resource usage for deployed mapping jobs that run in separate local processes. Detailed statistics appear for jobs that run longer than one minute.

1. In the Administrator tool, click the **Monitor** tab.
2. Click the **Execution Statistics** view.
3. In the Domain Navigator, expand a Data Integration Service.
4. Expand an application and select **Deployed Mapping Jobs**.

A list of mapping jobs appears in the contents panel.

5. Select a mapping job.

The details panel displays the Properties for the mapping job.

6. Click the **Detailed Statistics** view.

The **Detailed Statistics** view displays the throughput graph and resource usage graphs.

Optionally, you can complete the following tasks in the **Detailed Statistics** view:

Task	Description
Enlarge a graph	Move the cursor over a graph, and then click the magnifying glass icon.
Enlarge a section of an enlarged graph	Drag the cursor to select an area to enlarge.
Switch between rows and bytes in the throughput graph	Click the Bytes option or the Rows option.
Choose which statistics are plotted on the throughput graph	In the throughput field, select the sources and targets that you want to view.

Viewing Deployed Mapping Job Statistics with the REST Operations Hub Service

You can get monitoring statistics for deployed mapping jobs with the REST functionality.

Monitor Logical Data Objects

You can monitor logical data objects on the **Execution Statistics** view on the **Monitor** tab or the Monitoring tool.

You can view information about logical data objects included in an application. To monitor a logical data object, expand a Data Integration Service in the Navigator. Expand an application, and then select the **Logical Data Objects** folder. A list of logical data objects appears in the contents panel. The contents panel shows properties about each logical data object.

Select a logical data object in the contents panel to download the logs for a data object.

When you select a logical data object in the contents panel, the details panel shows the following views:

- **Properties** view
- **Cache Refresh Runs** view

Properties View for a Logical Data Object

The **Properties** view shows general properties and run-time statistics about the selected object.

You can view properties such as the data object name, logical data object model, folder path, cache state, and last cache refresh information.

Cache Refresh Runs View for a Logical Data Object

The **Cache Refresh Runs** view shows cache refresh details about the selected logical data object.

The **Cache Refresh Runs** view shows cache refresh details such as the cache run ID, request count, and row count.

Viewing Logs for Data Object Cache Refresh Runs

You can download the logs for data object cache refresh runs to view the cache refresh run details.

1. Click the **Execution Statistics** view.
2. In the Domain Navigator, expand a Data Integration Service.
3. Expand an application and select **Logical Data Objects**.
The contents panel displays a list of logical data objects.
4. Select a logical data object.
Details about the selected data object appear in the details panel.
5. Select the **Cache Refresh Runs** view.
6. Click **View Logs for Selected Object**.

Monitor SQL Data Services

You can monitor SQL data services on the **Execution Statistics** view on the **Monitor** tab. An SQL data service is a virtual database that you can query. It contains a schema and other objects that represent underlying physical data.

You can view information about the SQL data services included in an application. To monitor an SQL data service, expand an application in the Navigator and select the **SQL Data Services** folder. A list of SQL data services appears in the contents panel. The contents panel shows properties about each SQL data service, such as the name, description, and state.

When you select an SQL data service in the contents panel, the contents panel shows the following views:

- **Properties** view
- **Connections** view
- **Requests** view
- **Virtual Tables** view
- **Reports** view

Properties View for an SQL Data Service

The **Properties** view shows general properties and run-time statistics for an SQL data service.

When you select an SQL data service in the contents panel of the **Properties** view, you can view the general properties and run-time statistics.

General Properties for an SQL Data Service

You can view general properties, such as the SQL data service name and the description.

Statistics for an SQL Data Service

You can view run-time statistics about connections and requests for the SQL data service. Sample statistics include the number of connections to the SQL data service, the number of requests, and the number of aborted connections.

Connections View for an SQL Data Service

The **Connections** view displays properties about connections from third-party clients. The view shows properties such as the connection ID, state of the connection, connect time, elapsed time, and disconnect time.

When you select a connection in the contents panel, you can abort the connection or access the **Properties** view and **Requests** view in the details panel.

Properties View

The **Properties** view in the details panel shows the user who is using the connection, the state of the connection, and the connect time.

Requests View

The **Requests** view in the details panel shows information about the requests for the SQL connection. Each connection can have more than one request. The view shows request properties such as request ID, user name, state of the request, start time, elapsed time, and end time.

Aborting a Connection

You can abort a connection to prevent it from sending more requests to the SQL data service.

1. Click the **Execution Statistics** view.
2. In the Domain Navigator, expand a Data Integration Service.
3. Expand an application and select **SQL Data Services**.
The contents panel lists the SQL data services in the application.
4. Select an SQL data service.
The contents panel displays multiple views for the SQL data service.
5. Click the **Connections** view.
The contents panel lists connections to the SQL data service.
6. Select a connection.
7. Click **Actions > Abort Selected Connection**.

Requests View for an SQL Data Service

The **Requests** view displays properties for requests for each SQL connection.

The **Requests** view shows properties about the requests for the SQL connection. Each connection can have more than one request. The view shows request properties such as request ID, connection ID, user name, state of the request, start time, elapsed time, and end time.

Select a request in the contents panel to view additional information about the request in the details panel.

Aborting an SQL Data Service Connection Request

You can abort an SQL Data Service connection request. You might want to abort a connection request that hangs or that is taking an excessive amount of time to complete.

1. Click the **Execution Statistics** view.
2. In the Domain Navigator, expand a Data Integration Service.
3. Expand an application and select **SQL Data Services**.
The contents panel displays a list of SQL data services.
4. Select an SQL data service.
5. Click the **Requests** view.
A list of connection requests for the SQL data service appears.
6. Select a request row.
7. Click **Actions > Abort Selected Request**.

Viewing Logs for an SQL Data Service Request

You can download the logs for an SQL data service request to view the request details.

1. Click the **Execution Statistics** view.
2. In the Domain Navigator, expand a Data Integration Service.
3. Expand an application and select **SQL Data Services**.
The contents panel displays a list of SQL data services.
4. Select an SQL data service.

5. Click the **Requests** view.
A list of requests for the SQL data service appears.
6. Select a request row.
7. Click **Actions > View Logs for Selected Object**.

Virtual Tables View for an SQL Data Service

The **Virtual Tables** view displays properties about the virtual tables in the SQL data service.

The view shows properties about the virtual tables, such as the name and description. When you select a virtual table in the contents panel, you can view the **Properties** view and **Cache Refresh Runs** view in the details panel.

Properties View

The **Properties** view displays general information and run-time statistics about the selected virtual table. General properties include the virtual table name and the schema name. Monitoring statistics include the number of request, the number of rows cached, and the last cache refresh time.

Cache Refresh Runs View

The **Cache Refresh Runs** view displays cache information for the selected virtual table. The view includes the cache run ID, the request count, row count, and the cache hit rate. The cache hit rate is the total number of requests on the cache divided by the total number of requests for the data object.

Viewing Logs for an SQL Data Service Table Cache

You can download the logs for an SQL data service table cache to view the table cache details.

1. Click the **Execution Statistics** view.
2. In the Domain Navigator, expand a Data Integration Service.
3. Expand an application and select **SQL Data Services**.
The contents panel displays a list of SQL data services.
4. Select an SQL data service.
5. Click the **Virtual Tables** view.
A list of virtual tables for the SQL data service appears.
6. Select a table row.
Details about the selected table appear in the details panel.
7. Select the **Cache Refresh Runs** view.
8. Click **View Logs for Selected Object**.

Reports View for an SQL Data Service

The **Reports** view shows monitoring reports about the selected SQL data service.

When you monitor an SQL data service, the **Reports** view shows reports about the SQL data service. For example, you can view the Most Active SQL Connections report to determine the SQL connections that received the most connection requests during a specific time period.

Monitor Web Services

You can monitor web services on the **Execution Statistics** view on the **Monitor** tab. Web services are business functions that operate over the Web. They describe a collection of operations that are network accessible through standardized XML messaging.

You can view information about web services included in an application. To monitor a web service, expand an application in the Navigator and select the **Web Services** folder. A list of web services appears in the contents panel. The contents panel shows properties about each web service, such as the name, description, and state of each web service.

When you select the link for a web service in the contents panel, the contents panel shows the following views:

- **Properties** view
- **Reports** view
- **Operations** view
- **Requests** view

Properties View for a Web Service

The **Properties** view shows general properties and run-time statistics for a web service.

When you select a web service in the contents panel of the **Properties** view, you can view the general properties and monitoring statistics.

General Properties for a Web Service

You can view general properties about the web service, such as the name and type of object.

Statistics for a Web Service

You can view run-time statistics about web service requests during a specific time period. The **Statistics** section shows the number of completed, failed, and total web service requests.

Reports View for a Web Service

The **Reports** view shows monitoring reports about the selected web service.

When you monitor a web service, the **Reports** view shows reports about the web service. For example, you can view the Most Active WebService Client IP report to determine the IP addresses that received the most number of web service requests during a specific time period.

Operations View for a REST or SOAP Web Service

The **Operations** view shows the name and description of each operation or resource included in the web service. The view also displays properties, requests, and reports about each operation.

When you select a web service operation in the contents panel, the details panel shows the **Properties** view, **Requests** view, and **Reports** view.

Properties View

The **Properties** view shows general properties and statistics about the selected web service operation or resource. General properties include the operation or resource name and type of object. The view also

shows statistics about the web service operation during a particular time period. Statistics include the number of completed, failed, and total web service requests.

Requests View

The **Requests** view shows properties about each web service operation, such as request ID, user name, state, start time, elapsed time, and end time. You can filter the list of requests. You can also view logs for the selected web service request.

Reports View for a SOAP Web Service

The **Reports** view shows reports about SOAP web service operations.

Requests View for a Web Service

The **Requests** view shows properties about each web service request, such as request ID, user name, state, start time, elapsed time, and end time. You can filter the list of requests.

When you select a web service request in the contents panel, you can view logs about the request in the details panel. The details panel shows general properties and statistics about the selected web service request. Statistics include the number of completed, failed, and total web service requests.

You can also abort a web service request from the **Requests** view. To abort a web service request, select the workflow request and click **Actions > Abort Selected Request** in the contents panel.

Monitor Workflows

You can monitor workflows on the **Execution Statistics** view on the **Monitor** tab.

You can view information about workflow instances that run from a workflow in a deployed application. To monitor a workflow, expand an application in the Navigator and select the **Workflows** folder. A list of workflow instances appears in the contents panel. The contents panel shows properties about each workflow instance, such as the name, state, start time, and recovery properties of each workflow instance. If the workflow instance ran on a grid, the contents panel shows the node that ran each mapping in the workflow instance.

Select a workflow instance in the contents panel to perform the following tasks:

- View logs for the workflow instance.
- View the context of the workflow instance to view other workflow instances that started around the same time as the selected workflow instance.
- Cancel or abort the workflow instance.
- Recover the interrupted workflow instance.

Expand a workflow instance to view properties about the workflow objects.

Workflow Graph

You can view the details of a workflow that you run in the Monitoring tool in a graphical form.

After you run a workflow, you can see the graphical view of the workflow in the Monitoring tool. In the workflow graph, you can see the sequential run of the mapping tasks in the workflow. The workflow graph enables you to view the failure points in a workflow at a glance.

In the workflow graph, you can view the following details of a workflow:

- Mapping tasks in the workflow
- Task details
- Recovery details

You can perform the following tasks from the workflow graph:

- Abort a running workflow
- Cancel a running workflow
- Recover a failed workflow
- View the workflow logs

Viewing a Workflow Graph

You can view a workflow graph that shows the sequential run of the mapping tasks in the workflow.

1. Click the **Execution Statistics** view.
2. In the Domain Navigator, expand an application.
3. Select the **Workflows** folder.
A list of workflows appears in the contents panel.
4. Select the workflow that you want to view.
5. Click **Actions > View Workflow Graph**.
The workflow graph appears in a new window.

View Workflow Objects

When you expand a workflow instance in the contents panel, you can view properties about workflow objects, such as the name, state, start time, and elapsed time for the object.

Workflow objects include events, tasks, and gateways. When you monitor workflows, you can monitor the tasks that run in a workflow instance. The **Monitor** tab does not display information about events or gateways in the workflow instance.

If an expression in a conditional sequence flow evaluates to false, the Data Integration Service does not run the next object or any of the subsequent objects in that branch. The **Monitor** tab does not list objects that do not run in the workflow instance. When a workflow instance includes objects that do not run, the instance can still successfully complete.

You can expand a task in the contents panel to view information about the work item run by the task. For example, if the workflow contains a mapping task, you can view throughput and resource usage statistics for the mapping run.

Viewing Summary Statistics for Workflow Objects

You can view throughput and resource usage statistics for mapping objects in workflows that run in separate local processes.

1. Click the **Execution Statistics** view.
2. In the Domain Navigator, expand a Data Integration Service.
3. Expand an application and select the **Workflows** folder.
A list of workflows appears in the contents panel.

4. Expand a workflow that contains a mapping object.
5. Expand the mapping task and select the mapping.
6. In the details panel, click the **Summary Statistics** view.

The **Summary Statistics** view displays throughput and resource usage statistics for the source and target.

Optionally, you can sort the statistics in ascending or descending order. Click a column header to sort the column in ascending order. Click the column header again to sort the column in descending order.

Viewing Detailed Statistics for Workflow Objects

You can view graphs of the throughput and resource usage for mapping objects in workflows that run in separate local processes. Detailed statistics appear for jobs that run longer than one minute.

1. Click the **Execution Statistics** view.
2. In the Domain Navigator, expand a Data Integration Service.
3. Expand an application and select the **Workflows** folder.
A list of workflows appears in the contents panel.
4. Expand a workflow that contains a mapping object.
5. Expand the mapping task and select the mapping.
6. Click the **Detailed Statistics** view in the details panel.

The **Detailed Statistics** view displays the throughput graph and resource usage graphs.

Optionally, you can complete the following tasks in the **Detailed Statistics** view:

Task	Description
Enlarge a graph	Move the cursor over a graph, and then click the magnifying glass icon.
Enlarge a section of an enlarged graph	Drag the cursor to select an area to enlarge.
Switch between rows and bytes in the throughput graph	Click the Bytes option or the Rows option.
Choose which statistics are plotted on the throughput graph	In the throughput field, select the sources and targets that you want to view.

Workflow States

When you monitor a workflow instance, you can view the state of the workflow instance. If a workflow instance recovers after a task is interrupted, the Monitor adds an entry for the task instance that runs in the recovered workflow.

A workflow instance can have one of the following states:

Aborted

A workflow instance aborts when you choose to abort the workflow instance from the Monitoring tool or using the `infacmd wfs abortWorkflow` command. You can also choose to abort a running workflow instance when you stop the application that contains the workflow or when you disable the workflow in the application.

Note: A workflow instance also aborts if the active sequence flow in the workflow reaches a Terminate event.

Canceled

You choose to cancel the workflow instance from the **Monitor** tab or by using the `infacmd wfs cancelWorkflow` command.

The workflow can also enter a canceled state if the Data Integration Service shuts down unexpectedly. If the workflow is not configured for automatic recovery, the service process changes the workflow instance state to Canceled when the service process restarts. Before the Data Integration Service restarts, the workflow state and the task state appear as Running, although the workflow and the task are no longer running. If the workflow is configured for automatic recovery, the service process recovers the workflow instance and reruns the interrupted task when the service process restarts. The service process sets the workflow instance state to Running.

Completed

The Data Integration Service successfully completes the workflow instance. A completed workflow instance might indicate that all tasks, gateways, and sequence flow evaluations either successfully completed or were in a branch that did not run.

A workflow can also enter a Completed state if a Command, Mapping, Notification, or Human task encounters a recoverable error or nonrecoverable error. When the task encounters the error, the Data Integration Service fails the task. The Data Integration Service runs subsequent workflow objects if expressions in the conditional sequence flows evaluate to true or if the sequence flows do not include conditions. If the workflow instance finishes running without another interruption, the Data Integration Service updates the workflow state to Completed.

When the task fails, the Data Integration Service continues to run additional objects in the workflow instance if expressions in the conditional sequence flows evaluate to true or if the sequence flows do not include conditions. If the workflow instance finishes running without another interruption, the Data Integration Service updates the workflow state to completed. A completed workflow instance can contain both failed and completed tasks.

Failed

A workflow instance fails when a workflow error occurs. Workflow errors can occur when the Data Integration Service reads the parameter file at the start of the workflow run, copies workflow parameter and variable values to task input, or evaluates expressions in conditional sequence flows. In addition, a workflow error occurs if an Assignment task or a gateway fails.

When a workflow error occurs, the Data Integration Service stops processing additional objects and fails the workflow instance immediately. Workflow errors are nonrecoverable.

Running

The Data Integration Service is running the workflow instance.

Workflow Object States

Workflows include tasks and gateways. When you monitor a workflow instance, you can view the state of the tasks that run in the workflow instance.

Tasks can have one of the following states:

Aborted

A task aborts in the following situations:

- The task encounters a nonrecoverable error.
- You abort the workflow instance.

When you abort the workflow instance, the Data Integration Service first aborts the task and then aborts the workflow instance.

If you choose to abort the workflow instance while an Assignment task is running, the Data Integration Service completes running the task. The Data Integration Service then aborts the workflow instance and does not start running other objects.

Completed

The Data Integration Service successfully completes the task.

Failed

A task fails in the following situations:

- Any task in a workflow not enabled for recovery encounters any type of error.
- An Assignment task in a workflow enabled for recovery encounters any type of error.
- A Command, Mapping, Notification, or Human task with a restart recovery strategy in a workflow enabled for recovery encounters a non recoverable error.
- A Mapping task with a skip recovery strategy in a workflow enabled for recovery encounters any type of error.

Note: A workflow can complete if a task fails. The Data Integration Service runs subsequent workflow objects if expressions in the conditional sequence flows evaluate to true or if the sequence flows do not include conditions. If the workflow instance finishes running without another interruption, the Data Integration Service updates the workflow state to Completed.

Running

The Data Integration Service is running the task.

Mapping Task Work Item States

When you expand a Mapping task, you can view the state of the mapping run. When you expand a restarted Mapping task, you can view the mapping jobs run for each recovery attempt of the workflow instance. If a workflow instance recovers after a Mapping task is interrupted, the Monitor adds an entry for the task instance that runs in the recovered workflow.

You can also view the state of the mapping run from the workflow graph of the workflow that contains the mapping task.

Mappings run by a Mapping task can have one of the following states:

Aborted

The Mapping task aborts while the mapping is running because you choose to abort the workflow instance.

Completed

The Data Integration Service successfully completes the mapping.

Failed

The mapping encounters an error. The mapping and the Mapping task appear as Failed in the Monitor. The states do not depend on the Mapping task recovery strategy.

Running

The Data Integration Service is running the mapping.

Canceling or Aborting a Workflow

You can cancel or abort a workflow instance at any time. You might want to cancel or abort a workflow instance that stops responding or that is taking an excessive amount of time to complete.

When you cancel a workflow instance, the Data Integration Service finishes processing any running task and then stops processing the workflow instance. The service does not start running any subsequent workflow objects.

When you abort a workflow instance, the Data Integration Service attempts to kill the process on any running task. If an Assignment task or a gateway is running, the Data Integration Service completes the task or gateway. After the task aborts or completes, the service aborts the workflow instance. The service does not start running any subsequent workflow objects.

You can also cancel or abort a workflow from the workflow graph.

1. Click the **Execution Statistics** view.
2. In the Navigator, expand a Data Integration Service.
3. Expand an application and select **Workflows**.
A list of workflow instances appears in the contents panel.
4. Select a workflow instance.
5. Click **Actions > Cancel Selected Workflow** or **Actions > Abort Selected Workflow**.

Workflow Recovery

Workflow recovery is the completion of a workflow instance from the point of interruption.

When a workflow is enabled for recovery, you can recover a workflow instance if a task encounters a recoverable error, if you cancel the workflow instance, or if the Data Integration Service process shuts down unexpectedly.

View the workflow log to identify the cause of the interruption. After fixing any recoverable errors, you can recover the interrupted workflow instance if it is enabled for recovery.

You cannot change a workflow definition between the interrupted run and the recovery run. If a workflow instance has a recoverable state and you change the workflow metadata in the Developer tool and redeploy the application that contains the workflow, then the workflow instance is no longer recoverable.

The Data Integration Service tries to recover the previous workflow state if the service restarts after an unexpected shutdown. By default, the Data Integration Service does not recover a workflow instance that stopped during a Command task, Mapping task, or Notification task. In addition, the Data Integration Service cannot recover a workflow instance by default if you cancel the workflow instance or cancel a running task in the workflow instance. You can configure the recovery options on the workflow to enable the Data Integration Service to recover a workflow instance in such cases.

When you configure the workflow options, you can configure the workflow for manual recovery or automatic recovery. If you configure automatic recovery, the Data Integration Service restarts the workflow from the point of interruption without any human interaction. If you configure manual recovery, you can restart the workflow.

When a workflow instance recovers or when you recover a workflow instance, the Data Integration Service restarts the task. The service continues processing the subsequent workflow objects. If a workflow instance recovers after a task is interrupted, the Monitor adds an entry for the task instance that runs in the recovered

workflow. For example, if a workflow recovers three times and restarts a Mapping task each time, the Monitor contains three entries for the Mapping task.

Recovery Properties

The read-only recovery properties display for each workflow instance. You configure the recovery properties for the workflow definition in the Developer tool. You cannot change the values of the properties for the workflow instance.

The following table describes the read-only recovery properties for a workflow instance:

Property	Description
Recovery Enabled	Indicates that the workflow is enabled for recovery.
Automatically Recover Workflows	Indicates that the Data Integration Service process tries to automatically recover workflow instances that were interrupted. The workflow recovery starts after the Data Integration Service process restarts.

Recovering a Workflow

You can recover interrupted workflow instances that are enabled for recovery.

1. Click the **Execution Statistics** view.
2. In the Domain Navigator, expand a Data Integration Service.
3. Expand an application and select **Workflows**.
A list of workflow instances appears in the contents panel.
4. Select the interrupted workflow instance that you want to recover.
5. Click **Actions > Recover Selected Workflow**.
Monitor the state of the workflow recovery run in the contents panel.

Workflow Logs

The Data Integration Service generates log events when you run a workflow. Log events include information about workflow errors, task progress, and the setting of workflow variables. Log events also include the analyses of the links that the Data Integration Service evaluates in a sequence flow.

If a workflow instance includes a Mapping task, the Data Integration Service generates a separate log file for the mapping. The mapping log file includes any errors encountered during the mapping run and load summary and transformation statistics.

You can view the workflow and mapping logs from the Monitor tab.

When you recover an interrupted workflow instance, the Data Integration Service appends log events to the current workflow log. When the recovered workflow instance includes a Mapping task that is restarted, the Data Integration Service creates a mapping log.

If the workflow runs on a grid, the recovery of the workflow instance might run on a different node than the original workflow instance run. If the recovery runs on a different node and the log directory is not in a shared location, the Data Integration Service creates a log file with the same name on the current node.

Workflow Log Information

The information in the workflow log file represents the sequence of events that occur when the workflow runs.

The Data Integration Service writes information to the workflow log when the following types of event occur:

- The Data Integration Service starts to run a task or another object in the workflow.
- A task or another object in the workflow is in progress.
- The Data Integration Service finishes running a task or another object in the workflow.
- The Data Integration Service sets or updates a workflow variable.
- The Data Integration Service evaluates the links in a sequence flow and identifies the correct path for the workflow process.
- The workflow encounters a workflow error.

Viewing Logs for a Workflow

You can download the log for a workflow instance to view the workflow instance details.

1. In the Administrator tool, click the **Monitor** tab.
2. Click the **Execution Statistics** view.
3. In the Domain Navigator, expand a Data Integration Service.
4. Expand an application and select **Workflows**.
A list of workflow instances appears in the contents panel.
5. Select a workflow instance.
6. Click **Actions > View Logs for Selected Object**.
A dialog box appears with the option to open or save the log file.

Viewing Logs for a Mapping Run in a Workflow

You can download the log for a mapping run in a workflow to view the mapping details.

1. Click the **Execution Statistics** view.
2. In the Domain Navigator, expand a Data Integration Service.
3. Expand an application and select **Workflows**.
A list of workflow instances appears in the contents panel.
4. Expand a workflow instance.
5. Expand a Mapping task, and then select the mapping run by the task.
6. Click **Actions > View Logs for Selected Object**.
A dialog box appears with the option to open or save the log file.

Job Status After Application Service Restart or Failover

If the monitoring Model Repository Service restarts or fails over while the Data Integration Service is running jobs, the Monitoring tool provides the latest known status of all jobs after the monitoring Model Repository Service becomes available.

If the status of a job is not known, the Monitoring tool reports the status as UNKNOWN. When the monitoring Model Repository Service restarts or fails over to a backup node, it updates the Monitoring tool with the latest status of each job if the job is still running. The status of a job that completes before the monitoring Model Repository Service becomes available remains as UNKNOWN.

For example, a Data Integration Service runs on a grid. Two mappings are running when the monitoring Model Repository Service fails. The Monitoring tool does not have the latest status of these mappings. One mapping completes successfully before the monitoring Model Repository Service becomes available. The other mapping continues to run after the monitoring Model repository Service becomes available. The Monitoring tool reports the status of the first mapping as UNKNOWN. It shows the status of the second mapping as RUNNING.

Monitoring a Folder of Objects

You can view properties and statistics about objects in a folder that appears in the Navigator of the **Execution Statistics** view. You can select one of the following folders: Jobs, Deployed Mapping Jobs, Logical Data Objects, SQL Data Services, Web Services, or Workflows.

You can apply a filter to limit the number of objects that appear in the contents panel. You can create custom filters based on a time range. Custom filters allow you to select particular dates and times for job start times, end times, and elapsed times. Custom filters also allow you to filter results based on multiple filter criteria.

1. In the Administrator tool, click the **Monitor** tab.
2. Click the **Execution Statistics** view.
3. In the Domain Navigator, select the folder.
The contents panel shows a list of objects contained in the folder.
4. Right-click the header of the table to add or remove columns.
5. Select **Receive New Notifications** to dynamically display new jobs, operations, requests, or workflows in the **Monitor** tab.
6. Enter filter criteria to reduce the number of objects that appear in the contents panel.
7. Select the object in the contents panel to view details about the object in the details panel.
The details panel shows more information about the object selected in the contents panel.
8. To view jobs that started around the same time as the selected job, click **Actions > View Context**.
The selected job and other jobs that started around the same time appear in the **Context View** tab. You can also view the context of connections, deployed mappings, requests, and workflows.
9. Click the **Close** button to close the **Context View** tab.

Viewing the Context of an Object

View the context of an object to view other objects of the same type that started around the same time as the selected object. You might view the context of an object to troubleshoot a problem or to get a high-level understanding of what is happening at a particular period of time. You can view the context of jobs, deployed mappings, connections, requests, and workflows.

For example, you notice that your deployed mapping failed. When you view the context of the deployed mapping, an unfiltered list of deployed mappings appears in a separate working view, showing you all deployed mappings that started around the same time as your deployed mapping. You notice that the other deployed mappings also failed. You determine that the cause of the problem is that the Data Integration Service was unavailable.

1. In the Administrator tool, click the **Monitor** tab.
2. Click the **Execution Statistics** view.
3. In the Domain Navigator, expand a Data Integration Service and select the category of objects.
For example, select **Jobs**.
4. In the contents panel, select the object for which you want to view the context.
For example, select a job.
5. Click **Actions > View Context**.

Configuring the Date and Time Custom Filter

You can apply a custom filter on a Start Time or End Time column in the contents panel of the **Monitor** tab to filter results.

1. Select Custom as the filter option for the Start Time or End Time column.
The **Custom Filter: Date and Time** dialog box appears.
2. Enter the date range using the specified date and time formats.
3. Click **OK**.

Configuring the Elapsed Time Custom Filter

You can apply a custom filter on an Elapsed Time column in the contents panel of the **Monitor** tab to filter results.

1. Select Custom as the filter option for the Elapsed Time column.
The **Custom Filter: Elapsed Time** dialog box appears.
2. Enter the time range.
3. Click **OK**.

Configuring the Multi-Select Custom Filter

You can apply a custom filter on columns in the contents panel of the **Monitor** tab to filter results based on multiple selections.

1. Select Custom as the filter option for the column.
The **Custom Filter: Multi-Select** dialog box appears.
2. Select one or more filters.

3. Click **OK**.

CHAPTER 14

Log Management

This chapter includes the following topics:

- [Log Management Overview, 262](#)
- [Log Manager Architecture, 263](#)
- [Log Location, 265](#)
- [System Logs, 265](#)
- [Log Management Configuration, 266](#)
- [Using the Logs Tab, 267](#)
- [Log Events, 272](#)
- [Log Aggregator, 278](#)
- [Mapping Task Logs, 279](#)

Log Management Overview

The Service Manager accumulates log events for the domain, application services, users, and PowerCenter sessions and workflows. To perform the logging function, the Service Manager runs a Log Manager and a Log Agent.

The Log Manager runs on the master gateway node. It collects and processes log events for Service Manager domain operations, application services, and user activity. The log events contain operational and error messages for a domain. The Service Manager and the application services send log events to the Log Manager. When the Log Manager receives log events, it generates log event files. You can view service log events in the Administrator tool based on criteria that you provide.

The Log Agent runs on all nodes in the domain. The Log Agent retrieves the workflow and session log events that the PowerCenter Integration Service writes and displays them in the Workflow Monitor. Workflow log events include information about workflow processing, workflow errors, and tasks that the PowerCenter Integration Service performs. Session log events include information about the tasks performed by the PowerCenter Integration Service, session errors, and load summary and transformation statistics for the session. You can view log events for the last workflow run with the Log Events window in the Workflow Monitor.

The Log Agent also collects and processes log events for jobs that the Data Integration Service runs. These include profile jobs, scorecard jobs, preview jobs, mapping jobs, and SQL data services. You can view log events for these jobs on the Monitoring tab.

Log event files are binary files that the Administrator tool Logs Viewer uses to display log events. When you view log events in the Administrator tool, the Log Manager uses the log event files to display the log events for the domain, application services, and user activity.

Domain logs include domain, application service, and user activity logs. You can view them in the Administrator tool. System logs are for use only by Informatica Support to address open support issues.

You can use the Administrator tool to perform the following tasks with the Log Manager:

- Configure the log location. Configure the node that runs the Log Manager, the directory path for log event files, purge options, and time zone for log events.
- Configure log management. Configure the Log Manager to purge logs or purge logs manually. Save log events to XML, text, or binary files. Configure the time zone for the time stamp in the log event files.
- View log events. View domain function, application service, and user activity log events on the Logs tab. Filter log events by domain, application service type, and user.

Log Manager Architecture

The Service Manager on the master gateway node controls the Log Manager. The Log Manager starts when you start the Informatica services. After the Log Manager starts, it listens for log events from the Service Manager and application services. When the Log Manager receives log events, it generates log event files.

The Log Manager creates the following types of log files:

- Log event files. Stores log events in binary format. The Log Manager creates log event files to display log events in the Logs tab. When you view events in the Administrator tool, the Log Manager retrieves the log events from the event nodes.

The Log Manager stores the files by date and by node. Set the directory path with the `infasetup tools defineDomain -ld` option.

- Guaranteed Message Delivery files. Stores domain, application service, and user activity log events. The Service Manager writes the log events to temporary Guaranteed Message Delivery files and sends the log events to the Log Manager.

If the Log Manager becomes unavailable, the Guaranteed Message Delivery files stay in the default log directory on the node where the service runs. By default, the directory path is

`<Informatica_installation_directory>/logs/<Node_Name>`. When the Log Manager becomes available, the Service Manager for the node reads the log events in the temporary files, sends the log events to the Log Manager, and deletes the temporary files.

PowerCenter Session and Workflow Log Events

PowerCenter session and workflow logs are stored in a separate location from the domain, application service, and user activity logs. The PowerCenter Integration Service writes session and workflow log events to binary files on the node where the PowerCenter Integration Service runs.

The Log Manager performs the following tasks to process PowerCenter session and workflow log events:

1. During a session or workflow, the PowerCenter Integration Service writes binary log files on the node. It sends information about the logs to the Log Manager.
2. The Log Manager stores information about workflow and session logs in the domain database. The domain database stores information such as the path to the log file location, the node that contains the log, and the PowerCenter Integration Service that created the log.

3. When you view a session or workflow in the Log Events window of the Workflow Monitor, the Log Manager retrieves the information from the domain database. The Log Manager uses the information to determine the location of the logs.
4. The Log Manager dispatches a Log Agent to retrieve the log events on each node to display in the Log Events window.

Data Integration Service Job Log Events

Logs for the jobs that the Data Integration Service runs are stored in a separate location from the domain, application service, and user activity logs. The Data Integration Service writes job log events to text files on the node where it runs.

The Data Integration Service and the Log Manager perform the following tasks to process job log events for the Data Integration Service:

1. When the Data Integration Service runs a job, it writes log events to text files on the node. The Data Integration Service sends information about the logs to the Log Manager.
2. The Log Manager stores log information in the Model repository database. The Model repository database stores information such as the path to the log file location, the node that contains the log, and the Data Integration Service that created the log.
3. When you view a job log on the Monitor tab of the Administrator tool, the Log Manager retrieves the information from the Model repository database. The Log Manager uses the information to determine the location of the logs.
4. The Log manager dispatches a Log Agent to retrieve the log events on each node to display the log.

Log Manager Recovery

When a service generates log events, it sends them to the Log Manager on the master gateway node. When you have the high availability option and the master gateway node becomes unavailable, the application services send log events to the Log Manager on a new master gateway node.

The Service Manager, the application services, and the Log Manager perform the following tasks:

1. An application service process writes log events to a Guaranteed Message Delivery file.
2. The application service process sends the log events to the Service Manager on the gateway node for the domain.
3. The Log Manager processes the log events and writes log event files. The application service process deletes the temporary file.
4. If the Log Manager is unavailable, the Guaranteed Message Delivery files stay on the node running the service process. The Service Manager for the node sends the log events in the Guaranteed Message Delivery files when the Log Manager becomes available, and the Log Manager writes log event files.

Troubleshooting the Log Manager

Domain and application services write log events to Service Manager log files when the Log Manager cannot process log events. The Service Manager log files are located in the default logs directory. The Service Manager log files include `catalina.out`, `localhost_<date>.txt`, and `node.log`. Services write log events to different log files depending on the type of error.

Use the Service Manager log files to troubleshoot issues when the Log Manager cannot process log events. You will also need to use these files to troubleshoot issues when you contact Informatica Global Customer Support.

Note: You can troubleshoot an Informatica installation by reviewing the log files generated during installation. You can use the installation summary log file to find out which components failed during installation.

Log Location

The Service Manager on the master gateway node writes log event files to the log file directory. When you configure a node to serve as a gateway, you must configure the directory where the Service Manager on this node writes the log event files. Each gateway node must have access to the directory path.

You configure the log location in the Properties view for the domain. Configure a directory location that is accessible to the gateway node during installation or when you define the domain. Store the logs on a shared disk when you have more than one gateway node. If the Log Manager is unable to write to the directory path, it writes log events to node.log on the master gateway node.

When you configure the log location, the Administrator tool validates the directory as you update the configuration. If the directory is invalid, the update fails. The Log Manager verifies that the log directory has read/write permissions on startup. Log files might contain inconsistencies if the log directory is not shared in a highly available environment.

You can change the directory path for domain logs in the Administrator tool or with the log service directory parameter, `-ld`. You can use the `-ld` parameter with any of the following commands:

- `infacmd isp SwitchToGatewayNode`
- `infasetup DefineDomain`
- `infasetup DefineGatewayNode`
- `infasetup UpdateGatewayNode`

System Logs

System logs contain information that Informatica Support views to help solve issues that you raise with Support. Ordinarily, you have no need to view these logs.

By default, the directory path is `<Informatica_installation_directory>/logs/<Node_Name>/..`. You can change the default directory path for logs with the System Log Directory parameter, `-sld`. You can use the `-sld` parameter with any of the following commands:

- `infasetup DefineDomain`
- `infasetup DefineGatewayNode`
- `infasetup DefineWorkerNode`
- `infasetup UpdateGatewayNode`
- `infasetup UpdateWorkerNode`

When you create a custom location, you can use a local location or a location that all domain nodes share. The Service Manager adds the node name to the path and creates separate log directories for each node.

When you update the gateway node or worker node with a new default location for system logs, existing logs remain intact. The server creates future logs at the new location, and abandons logs at the old location.

If you specify a node name when you change the default path, the Service Manager adds it to the path. For example, if you specify C:/logs/node1/ as the system log directory, the Service Manager creates logs in directories under C:/logs/node1/node1/.

If you have multiple Informatica domains, you must configure a different directory path for the Log Manager in each domain. Multiple domains cannot use the same shared directory path.

Note: When you change the directory path, you must restart Informatica services on the node you changed.

Log Management Configuration

The Service Manager and the application services continually send log events to the Log Manager. As a result, the directory location for the logs can grow to contain a large number of log events.

You can purge logs events periodically to manage the amount of log events stored by the Log Manager. You can export logs before you purge them to keep a backup of the log events.

Purging Log Events

You can automatically or manually purge log events. The Service Manager purges log events from the log directory according to the purge properties you configure in the Log Management dialog box. You can manually purge log events to override the automatic purge properties.

Purging Log Events Automatically

The Service Manager purges log events from the log directory according to the purge properties.

When the number of days or the size of the log directory exceeds the limit, the Log Manager deletes the log event files, starting with the oldest log events. The Log Manager periodically verifies the purge options and purges log events. The Log Manager does not purge the current day log event files and folder.

The following table lists purge properties:

Option	Description
Preserve logs for number of days	Number of days to preserve logs. Default is 30.
Maximum size for logs in MB	Number of megabytes of disk space to store logs. Default is 200.

Note: The Log Manager does not purge PowerCenter session and workflow log files.

Purging Log Events Manually

You can purge log events for the domain, application services, or user activity. When you purge log events, the Log Manager removes the log event files from the log directory. The Log Manager does not remove log event files currently being written to the logs.

Optionally, you can use the *infacmd* PurgeLog command to purge log events.

The following table lists the purge log options:

Option	Description
Log Type	Type of log events to purge. You can purge domain, service, user activity or all log events.
Service Type	When you purge application service log events, you can purge log events for a particular application service type or all application service types.
Purge Entries	Date range of log events you want to purge. You can select the following options: <ul style="list-style-type: none">- All Entries. Purges all log events.- Before Date. Purges log events that occurred before this date. Use the yyyy-mm-dd format when you enter a date. Optionally, you can use the calendar to choose the date. To use the calendar, click the date field.

Time Zone

When the Log Manager creates log event files, it generates a time stamp based on the time zone for each log event. When the Log Manager creates log folders, it labels folders according to a time stamp. When you export or purge log event files, the Log Manager uses this property to calculate which log event files to purge or export. Set the time zone to the location of the machine that stores the log event files.

Verify that you do not lose log event files when you configure the time zone for the Log Manager. If the application service that sends log events to the Log Manager is in a different time zone than the master gateway node, you may lose log event files you did not intend to delete. Configure the same time zone for each gateway node.

Note: When you change the time zone, you must restart Informatica Services on the node that you changed.

Configuring Log Management Properties

Configure the log management properties in the **Log Management** dialog box in Informatica Administrator.

1. In the Administrator console, click the **Logs** tab.
2. Select **Log Actions > Log Management**.
3. Enter the number of days for the Log Manager to preserve log events.
4. Enter the maximum disk size for the directory that contains the log event files.
5. Enter the time zone in the following format:
GMT (+|-) <hours>:<minutes>
For example: GMT+08:00
6. Click **OK**.

Using the Logs Tab

You can view domain, application service, and user activity log events in the Logs tab of the Administrator tool. When you view log events in the Logs tab, the Log Manager displays the generated log event files in the log directory. When an error message appears in the Administrator tool, the error provides a link to the Logs tab.

You can use the Logs tab to perform the following tasks:

- View log events and the Administrator tool operational errors. View log events for the domain, an application service, or user activity.
- Filter log event results. After you display the log events, you can display log events that match filter criteria.
- Configure columns. Configure the columns you want the Logs tab to display.
- Save log events. You can save log events in XML, text, and binary format.
- Purge log events. You can manually purge log events.
- Copy log event rows. You can copy log event rows.

Viewing Log Events

To view log events in the Logs tab of the Administrator tool, select the Domain, Service, or User Activity view. Next, configure the filter options. You can filter log events based on attributes such as log type, domain function category, application service type, application service name, user, message code, activity code, timestamp, and severity level. The available options depend on whether you choose to view domain, application service, or user activity log events.

To view more information about a log event, click the log event in the search results.

On AIX and Linux, if the Log Manager receives an internal error message from the PowerCenter Integration Service, it writes a stack trace to the log event window.

You can view logs to get more information about errors that you receive while working in the Administrator tool.

1. In the Administrator Tool, click the Logs tab.
2. In the contents panel, select Domain, Service, or User Activity view.
3. Configure the filter criteria to view a specific type of log event.

The following table lists the query options:

Log Type	Option	Description
Domain	Category	Category of domain service you want to view.
Service	Service Type	Application service you want to view.
Service	Service Name	Name of the application service for which you want to view log events. You can choose a single application service name or all application services.
Domain, Service	Severity	The Log Manager returns log events with this severity level.
User Activity	User	User name for the Administrator tool user.
User Activity	Security Domain	Security domain to which the user belongs.

Log Type	Option	Description
Domain, Service, User Activity	Timestamp	Date range for the log events that you want to view. You can choose the following options: <ul style="list-style-type: none"> - Blank. View all log events. - Within Last Day - Within Last Month - Custom. Specify the start and end date. Default is Within Last Day.
Domain, Service	Thread	Filter criteria for text that appears in the thread data. You can use wildcards (*) in this text field.
Domain, Service	Message Code	Filter criteria for text that appears in the message code. You can also use wildcards (*) in this text field.
Domain, Service	Message	Filter criteria for text that appears in the message. You can also use wildcards (*) in this text field.
Domain, Service	Node	Name of the node for which you want to view log events.
Domain, Service	Process	Process identification number for the Windows or UNIX service process that generated the log event. You can use the process identification number to identify log events from a process when an application service runs multiple processes on the same node.
User Activity	Activity Code	Filter criteria for text that appears in the activity code. You can also use wildcards (*) in this text field.
User Activity	Activity	Filter criteria for text that appears in the activity. You can also use wildcards (*) in this text field.

4. Click the Filter button.

The Log Manager retrieves the log events and displays them in the Logs tab with the most recent log events first.

5. Click the Reset Filter button to view a different set of log events.

Tip: To search for logs related to an error or fatal log event, note the timestamp of the log event. Then, reset the filter and use a custom filter to search for log events during the timestamp of the event.

Configuring Log Columns

You can configure the Logs tab to display the following columns:

- Category
- Service Type
- Service Name
- Severity
- User
- Security Domain

- Timestamp
- Thread
- Message Code
- Message
- Node
- Process
- Activity Code
- Activity

Note: The columns appear based on the query options that you choose. For example, when you display a service type, the service name appears in the Logs tab.

1. In the Administrator Tool, click the **Logs** tab.
 2. Select the **Domain, Service, or User Activity** view.
 3. To add a column, right-click a column name, select **Columns**, and then the name of the column you want to add.
 4. To remove a column, right-click a column name, select **Columns**, and then clear the checkmark next to the name of the column you want to remove.
 5. To move a column, select the column name, and then drag it to the location where you want it to appear.
- The Log Manager updates the Logs tab columns with your selections.

Saving Log Events

You can save the log events that you filter and view in the Log Viewer. When you save log events, the Log Manager saves whatever logs that you are viewing based on the filter criteria. To save log events to a file, click Save Logs on the Log Actions menu.

The Log Manager does not delete the log events when you save them. The Administrator Tool prompts you to save or open the saved log events file.

Optionally, you can use the *infacmd* `isp GetLog` command to retrieve log events.

The format you choose to save log events to depends on how you plan to use the exported log events file:

- XML file. Use XML format if you want to analyze the log events in an external tool that uses XML or if you want to use XML tools, such as XSLT.
- Text file. Use a text file if you want to analyze the log events in a text editor.
- Binary file. Use binary format to back up the log events in binary format. You might need to use this format to send log events to Informatica Global Customer Support.

Exporting Log Events

You can export the log events to an XML, text, or binary file. To export log events to a file, click Export Logs on the Log Actions menu.

When you export log events, you can choose which logs you want to save. When you choose Service logs, you can export logs for a particular service type. You can choose the sort order of the log events in the export file.

The Log Manager does not delete the log events when you export them. The Administrator tool prompts you to save or open the exported log events file.

Optionally, you can use the *infacmd* `GetLog` command to retrieve log events.

The format you choose to export log events depends on how you plan to use the exported log events file:

- XML file. Use XML format if you want to analyze the log events in an external tool that uses XML or if you want to use XML tools, such as XSLT.
- Text file. Use a text file if you want to analyze the log events in a text editor.
- Binary file. Use binary format to back up the log events in binary format. You might need to use this format to send log events to Informatica Global Customer Support.

The following table describes the export log options for each log type:

Option	Log Type	Description
Type	Domain, Service, User Activity	Type of logs you want to export.
Service Type	Service	Type of application service for which to export log events. You can also export log events for all service types.
Export Entries	Domain, Service, User Activity	Date range of log events you want to export. You can select the following options: <ul style="list-style-type: none"> - All Entries. Exports all log events. - Before Date. Exports log events that occurred before this date. Use the yyyy-mm-dd format when you enter a date. Optionally, you can use the calendar to choose the date. To use the calendar, click the date field.
Export logs in descending chronological order	Domain, Service, User Activity	Exports log events starting with the most recent log events.

XML Format

When you export log events to an XML file, the Log Manager exports each log event as a separate element in the XML file. The following example shows an excerpt from a log events XML file:

```
<log xmlns:xsd="http://www.w3.org/2001/XMLSchema" xmlns:common="http://www.informatica.com/pcsf/common" xmlns:metadata="http://www.informatica.com/pcsf/metadata" xmlns:domainservice="http://www.informatica.com/pcsf/domainservice" xmlns:logservice="http://www.informatica.com/pcsf/logservice" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
  <logEvent xsi:type="logservice:LogEvent" objVersion="1.0.0" timestamp="1129098642698" severity="3" messageCode="AUTHEN_USER_LOGIN_SUCCEEDED" message="User Admin successfully logged in." user="Admin" stacktrace="" service="authenticationservice" serviceType="PCSF" clientNode="sapphire" pid="0" threadName="http-8080-Processor24" context="" />
  <logEvent xsi:type="logservice:LogEvent" objVersion="1.0.0" timestamp="1129098517000" severity="3" messageCode="LM_36854" message="Connected to node [garnet] on outbound connection [id = 2]." user="" stacktrace="" service="Copper" serviceType="IS" clientNode="sapphire" pid="4484" threadName="4528" context="" />
</log>
```

Text Format

When you export log events to a text file, the Log Manager exports the log events in Information and Content Exchange (ICE) Protocol. The following example shows an excerpt from a log events text file:

```
2006-02-27 12:29:41 : INFO : (2628 | 2768) : (IS | Copper) : sapphire : LM_36522 :
Started process [pid = 2852] for task instance Session task instance
[s_DP_m_DP_AP_T_DISTRIBUTORS4]:Executor - Master.
```

```
2006-02-27 12:29:41 : INFO : (2628 | 2760) : (IS | Copper) : sapphire : CMN_1053 :
Starting process [Session task instance [s_DP_m_DP_AP_T_DISTRIBUTORS4]:Executor -
Master].
2006-02-27 12:29:36 : INFO : (2628 | 2760) : (IS | Copper) : sapphire : LM_36522 :
Started process [pid = 2632] for task instance Session task instance
[s_DP_m_DP_AP_T_DISTRIBUTORS4]:Preparer.
2006-02-27 12:29:35 : INFO : (2628 | 2760) : (IS | Copper) : sapphire : CMN_1053 :
Starting process [Session task instance [s_DP_m_DP_AP_T_DISTRIBUTORS4]:Preparer].
```

Binary Format

When you export log events to a binary file, the Log Manager exports the log events to a file that Informatica Global Customer Support can import. You cannot view the file unless you convert it to text. You can use the *infacmd* ConvertLogFile command to convert binary log files to text files, XML files, or readable text on the screen.

Viewing Administrator Tool Log Errors

If you receive an error while starting, updating, or removing services in the Administrator tool, an error message in the contents panel of the service provides a link to the Logs tab. Click the link in the error message to access detail information about the error in the Logs tab.

Log Events

The Service Manager and application services send log events to the Log Manager. The Log Manager generates log events for each service type.

Log events include a timestamp, in milliseconds, and a thread name that identifies the event.

You can view the following log event types on the Logs tab:

- Domain log events. Log events generated from the Service Manager functions.
- Analyst Service log events. Log events about each Analyst Service running in the domain.
- Content Management Service log events. Log events about each Content Management Service running in the domain.
- Data Integration Service log events. Log events about each Data Integration Service running in the domain.
- Metadata Manager Service log events. Log events about each Metadata Manager Service running in the domain.
- Model Repository log events. Log events about each Model Repository Service running in the domain.
- PowerCenter Integration Service log events. Log events about each PowerCenter Integration Service running in the domain.
- PowerCenter Repository Service log events. Log events from each PowerCenter Repository Service running in the domain.
- Resource Manager Service log events. Log events about the Resource Manager Service running in the domain.
- SAP BW Service log events. Log events about the interaction between the PowerCenter and the SAP NetWeaver BI system.

- Web Services Hub log events. Log events about the interaction between applications and the Web Services Hub.
- User activity log events. Log events about domain and security management tasks that a user completes.

Log Event Components

The Log Manager uses a common format to store and display log events. You can use the components of the log events to troubleshoot Informatica.

Each log event contains the following components:

- Service type, category, or user. The Logs tab categorizes events by domain category, service type, or user. If you view application service logs, the Logs tab displays the application service names. When you view domain logs, the Logs tab displays the domain categories in the log. When you view user activity logs, the Logs tab displays the users in the log.
- Message or activity. Message or activity text for the log event. Use the message text to get more information about the log events for domain and application services. Use the activity text to get more information about log events for user activity. Some log events contain embedded log event in the message texts. For example, the following log events contains an embedded log event:

```
Client application [PmDTM], connection [59]: recv failed.
```

In this log event, the following log event is the embedded log event:

```
[PmDTM], connection [59]: recv failed.
```

When the Log Manager displays the log event, the Log Manager displays the severity level for the embedded log event.

- Security domain. When you view user activity logs, the Logs tab displays the security domain for each user.
- Message or activity code. Log event code. If the message type is error or fatal, click on the message code to open the Informatica Knowledge Base search for the message. You must configure the support portal credentials in the user account to do the search.
- Process. The process identification number for the Windows or UNIX service process that generated the log event. You can use the process identification number to identify log events from a process when an application service runs multiple processes on the same node.
- Node. Name of the node running the process that generated the log event.
- Thread. Identification number or name of a thread started by a service process.
- Time stamp. Date and time the log event occurred.
- Severity. The severity level for the log event. When you view log events, you can configure the Logs tab to display log events for a specific severity level.

Domain Log Events

Domain log events are log events generated from the domain functions the Service Manager performs.

Use the domain log events to view information about the domain and troubleshoot issues. You can use the domain log events to troubleshoot issues related to the startup and initialization of nodes and application services for the domain.

Domain log events include log events from the following functions:

- Authorization. Log events that occur when the Service Manager authorizes user requests for services. Requests can come from the Administrator tool.

- Container Management. Log events that occur when the Service Manager manages containers on nodes with the compute role.
- Domain Configuration. Log events that occur when the Service Manager manages the domain configuration metadata.
- Licensing. Log events that occur when the Service Manager registers license information.
- License Usage. Log events that occur when the Service Manager verifies license information from application services.
- Log Manager. Log events from the Log Manager. The Log Manager runs on the master gateway node. It collects and processes log events for Service Manager domain operations and application services.
- Log Agent. Log events from the Log Agent. The Log Agent runs on all nodes in the domain. It retrieves PowerCenter workflow and session log events to display in the Workflow Monitor.
- Monitoring. Log events about Domain Functions.
- Node Configuration. Log events that occur as the Service Manager manages node configuration metadata in the domain.
- User Management. Log events that occur when the Service Manager manages users, groups, roles, and privileges.
- Service Manager. Log events from the Service Manager and signal exceptions from DTM processes. The Service Manager manages all domain operations. If the error severity level of a node is set to Debug, when a service starts the log events include the environment variables used by the service.

Analyst Service Log Events

Analyst Service log events contain the following information:

- Managing projects. Log events about managing projects in the Informatica Analyst, such as creating objects, folders, and projects. Log events about creating profiles, scorecards, and reference tables.
- Running jobs. Log events about running profiles and scorecards. Logs about previewing data.
- User permissions. Log events about managing user permissions on projects.

Data Integration Service Log Events

Data Integration Service logs contain logs about the following events:

- Configuration. Log events about system or service configuration changes, application deployment or removal, and logs about the associated profiling warehouse.
- Data Integration Service processes. Log events about application deployment, data object cache refresh, and user requests to run mappings, jobs, or workflows.
- Service failures. Log events about failures that cause the Data Integration Service to be unavailable, such as Model repository connection failures or the service's failure to start.

Listener Service Log Events

The PowerExchange Listener logs contain information about the application service that manages the PowerExchange Listener.

The Listener Service logs contain the following information:

- Client communication. Log events for communication between a PowerCenter or PowerExchange client and a data source.

- Listener service. Log events about the Listener service, including configuring, enabling, and disabling the service.
- Listener service operations. Log events for operations such as managing bulk data movement and change data capture.

Logger Service Log Events

The PowerExchange Logger Service writes logs about the application service that manages the PowerExchange Logger.

The Logger Service logs contain the following information:

- Connections. Log events about connections between the Logger Service and the source databases.
- Logger service. Log events about the Logger Service, including configuring, enabling, and disabling the service.
- Logger service operations. Log events for operations such as capturing changed data and writing the data to PowerExchange Logger files.

Model Repository Service Log Events

Model repository Service log events contain the following information:

- Model repository connections. Log events for connections to the repository from Informatica Developer, Informatica Analyst, and the Data Integration Service.
- Model Repository Service. Log events about the Model Repository Service, including enabling, disabling, starting, and stopping the service.
- Repository operations. Log events for repository operations such as creating and deleting repository content, and adding deployed applications.
- User permissions. Log events about managing user permissions on the repository.

Metadata Manager Service Log Events

The Metadata Manager Service log events contain information about each Metadata Manager Service running in the domain.

Metadata Manager Service log events contain the following information:

- Repository operations. Log events for accessing metadata in the Metadata Manager repository.
- Configuration. Log events about the configuration of the Metadata Manager Service.
- Run-time processes. Log events for running a Metadata Manager Service, such as missing native library files.
- PowerCenter Integration Service log events. Session and workflow status for sessions and workflows that use a PowerCenter Integration Service process to load data to the Metadata Manager warehouse or to extract source metadata.

To view log events about how the PowerCenter Integration Service processes a PowerCenter workflow to load data into the Metadata Manager warehouse, you must view the session or workflow log.

PowerCenter Integration Service Log Events

The PowerCenter Integration Service log events contain information about each PowerCenter Integration Service running in the domain.

PowerCenter Integration Service log events contain the following information:

- PowerCenter Integration Service processes. Log events about the PowerCenter Integration Service processes, including service ports, code page, operating mode, service name, and the associated repository and PowerCenter Repository Service status.
- Licensing. Log events for license verification for the PowerCenter Integration Service by the Service Manager.

PowerCenter Repository Service Log Events

The PowerCenter Repository Service log events contain information about each PowerCenter Repository Service running in the domain.

PowerCenter Repository Service log events contain the following information:

- PowerCenter Repository connections. Log events for connections to the repository from PowerCenter client applications, including user name and the host name and port number for the client application.
- PowerCenter Repository objects. Log events for repository objects locked, fetched, inserted, or updated by the PowerCenter Repository Service.
- PowerCenter Repository Service processes. Log events about PowerCenter Repository Service processes, including starting and stopping the PowerCenter Repository Service and information about repository databases used by the PowerCenter Repository Service processes. Also includes repository operating mode, the nodes where the PowerCenter Repository Service process runs, initialization information, and internal functions used.
- Repository operations. Log events for repository operations, including creating, deleting, restoring, and upgrading repository content, copying repository contents, and registering and unregistering local repositories.
- Licensing. Log events about PowerCenter Repository Service license verification.

Resource Manager Service Log Events

Resource Manager Service log events contain the following information:

- Resource Manager Service. Log events about the Resource Manager Service, including enabling, disabling, starting, and stopping the service.
- Compute nodes. Log events about nodes with the compute role registering with the Resource Manager Service.

SAP BW Service Log Events

The SAP BW Service log events contain information about the interaction between PowerCenter and the SAP NetWeaver BI system.

SAP NetWeaver BI log events contain the following log events for an SAP BW Service:

- SAP NetWeaver BI system log events. Requests from the SAP NetWeaver BI system to start a workflow and status information from the ZPMSENDSTATUS ABAP program in the process chain.
- PowerCenter Integration Service log events. Session and workflow status for sessions and workflows that use a PowerCenter Integration Service process to load data to or extract data from SAP NetWeaver BI.

To view log events about how the PowerCenter Integration Service processes an SAP NetWeaver BI workflow, you must view the session or workflow log.

Scheduler Service Log Events

Scheduler Service logs contain information about the following events:

- Scheduler Service events. Log events about the Scheduler Service, including enabling, disabling, starting, and stopping the service.
- Scheduled object events. Log events about starting scheduled object runs.

Web Services Hub Log Events

The Web Services Hub log events contain information about the interaction between applications and the Web Services Hub.

Web Services Hub log events contain the following log events:

- Web Services processes. Log events about web service processes, including starting and stopping Web Services Hub, web services requests, the status of the requests, and error messages for web service calls. Log events include information about which service workflows are fetched from the repository.
- PowerCenter Integration Service log events. Workflow and session status for service workflows including invalid workflow errors.

User Activity Log Events

User activity log events describe all domain and security management tasks that a user completes.

Use the user activity log events to determine when a user created, updated, or removed services, nodes, users, groups, or roles.

The Service Manager writes user activity log events when the Service Manager needs to authorize a user to perform one of the following domain actions:

- Enable or disable a service process.
- Start, stop, enable, or disable a service.
- Add, update, or shut down a node.
- Modify the domain properties.
- Move a folder in the domain.

The Service Manager also writes user activity log events each time a user adds, updates, or removes a user, group, operating system profile, or role.

The user activity log displays information about the user who performed the security action.

The Service Manager writes a user activity log event each time a user account is locked or unlocked. The Service Manager also writes a user activity log event each time a user tries to log in to the domain with a client application.

The user activity logs also displays information about security audit trails and log events for changes to users, groups, and permissions.

To include security audit trails in the user activity log events, you must enable the SecurityAuditTrail property for the PowerCenter Repository Service in the Administrator tool.

When you import one or more repository objects, you can generate audit logs.

The audit logs contain the following information about the .xml file imported:

- Host name and IP address of the client machine from which the .xml file was imported

- Full local path of the .xml import file
- The file name
- The file size in bytes
- Logged in user name
- Number of objects imported
- Time stamp of the import operation

Log Aggregator

You can aggregate the log files of an application service that stops responding or shuts down unexpectedly. You might need to analyze multiple log files to figure out issues with an application service.

You can use the log aggregator to aggregate all the log files associated with an application service and compress required log files into a .zip file. You can download the .zip file and analyze the log files, or upload the .zip file to Informatica Global Customer Support for analysis.

You cannot store the history of aggregated logs. You must download or send the file to Informatica Global Customer Support after you aggregate the log files.

You can aggregate the hang and crash logs of the following application services:

- Analyst Service
- Data Integration Service
- Model Repository Service
- PowerCenter Integration Service
- PowerCenter Repository Service

In addition to the application service logs, the log aggregator captures debug information for the nodes in the domain. The log aggregator aggregates the log files of the associated application services when you aggregate the log files of an application service. For example, when you aggregate the log files of an Analyst Service, the log aggregator aggregates the log files of the Data Integration Service and the Model Repository Service associated with the Analyst Service.

The log collection directory in the master gateway node stores the application service logs when you aggregate the logs. All the node processes in the domain must have read/write access on the log collection directory. If the node processes cannot access the log collection directory, the aggregated logs do not appear in the aggregated logs listgrid. The core dump directory stores the core dump files of the nodes in the domain. Configure the log collection directory in the master gateway node and the core dump directory for each node in the domain.

When you process the aggregated logs you can choose the collectors from which you want to collect log information. The collectors are application services and nodes associated with the application service.

Aggregating Application Service Logs

You can aggregate log files associated with hang or crash scenarios of an application service.

1. Click the **Logs** tab in the Administrator tool.
2. Click the **Log Aggregator** tab.
3. Select the application service for which you want to aggregate the logs.

4. Select the scenario for which you want to aggregate the logs.
You can choose between application service crash and hang scenarios.
5. Select the time interval to aggregate the logs.
You can choose to aggregate the logs from the previous 6 hours to 3 days.
6. Click **Next**.
7. Select the collectors from which you want to aggregate the logs.
The log aggregator displays the log files and the collectors based on the node to which they belong.
8. Click **Finish**.
The list of logs associated with the scenario appears in the right pane. You can download the aggregated logs or send the logs to the Informatica Global Customer Support.

Processing Aggregated Application Service Logs

After you aggregate the application service logs, you must download the aggregated zip file or send the logs to Informatica Global Customer Support.

Aggregate the application service logs based on your requirement.

1. Select the logs that you want to process.
2. Click **Actions > Compress Logs**.
The **Compressed Scenario Output** dialog box appears.
3. On the **Compressed Output** tab, click **Download** to download the aggregated log files as a zip file.
4. Optionally, click the **Send to Support** tab.
5. Enter the user name, password, and the TFTP directory of the Informatica My Support Portal.
6. Click **Send** to send the aggregated log files to Informatica Global Customer Support.

Mapping Task Logs

You can view Mapping task logs to troubleshoot Mapping task problems or to see information about the mapping run.

The Data Integration service writes a new log file for each Mapping task run. The log file contains information about the events in the Mapping task. Log events are lines of text that contain a timestamp, thread identifier, a severity code, and the log message. The message can contain general information or it can contain an error message.

The following text shows the Mapping task log message format:

```
2015-02-20 12:49:24 <DTMLoggerThread_2> INFO: READER_1_1_1,   DBG_21430,   Reading
data from input source file [C:\Source\Logging_Source_1.txt]
2015-02-20 12:49:24 <DTMLoggerThread_2> INFO: READER_1_1_1,   BLKR_16019,   Read [200]
rows, read [0] error rows for source table [read_src2] instance name [read_src2]
2015-02-20 12:49:24 <DTMLoggerThread_2> INFO: LKPPD_2,   TE_7212,   Increasing [Data
Cache] size for transformation [Rel_Lookup] from [59652322] to [59654144].
2015-02-20 12:49:24 <DTMLoggerThread_2> INFO: READER_1_1_1,   BLKR_16008,   Reader run
completed.
2015-02-20 12:49:24 <DTMLoggerThread_2> INFO: WRITER_1_*_1,   WRT_8167,   Start
loading table [Router_Target_Default] at: Fri Feb 20 12:49:24 2015
```

When you set the tracing level to `verboseData`, the Mapping task log shows the parameters and parameter values for the mapping run.

The following text shows some Mapping task log messages that contain parameter values:

```
Integration Service will use override value [C:\Source] for parameter [ff_SrcDir] in
transformation [map_AllTx\read_src1].
Integration Service will use override value [8] for parameter [exp_Int] in
transformation [map_AllTx\Expression].
Integration Service will use override value [Mapping_New] for parameter [exp_String] in
transformation [map_AllTx\Expression].
Integration Service will use override value [C:\Source] for parameter [ldo_SrcDir] in
mapping \ mapplet [map_AllTx\DO_Lookup\DO_FF_REL_SRC_Read_Mapping].
```

After you run a mapping on the Spark engine on a Hadoop cluster, you can view the total number of cluster nodes used to execute the mapping in the mapping task log. On the Blaze engine, you can view the number of healthy cluster nodes used by the Grid Manager in the mapping task log.

CHAPTER 15

Domain Reports

This chapter includes the following topics:

- [Domain Reports Overview, 281](#)
- [License Management Report, 281](#)
- [Web Services Report, 288](#)

Domain Reports Overview

You can run the following domain reports from the Reports tab in the Administrator tool:

- **License Management Report.** Monitors the number of software options purchased for a license and the number of times a license exceeds usage limits. The License Management Report displays the license usage information such as CPU and repository usage and the node configuration details.
- **Web Services Report.** Monitors activities of the web services running on a Web Services Hub. The Web Services Report displays run-time information such as the number of successful or failed requests and average service time. You can also view historical statistics for a specific period of time.

Note: If the master gateway node runs on a UNIX machine and the UNIX machine does not have a graphics display server, you must install X Virtual Frame Buffer on the UNIX machine to view the report charts in the License Report or the Web Services Report. If you have multiple gateway nodes running on UNIX machines, install X Virtual Frame Buffer on each UNIX machine.

License Management Report

You can monitor the list of software options purchased with a license and the number of times a license exceeds usage limits. The License Management Report displays the general properties, CPU and repository usage, user details, hardware and node configuration details, and the options purchased for each license.

You can save the License Management Report as a PDF on your local machine. You can also email a PDF version of the report to someone.

Run the License Management Report to monitor the following license usage information:

- **Licensing details.** Shows general properties for every license assigned in the domain.
- **CPU usage.** Shows the number of logical CPUs used to run application services in the domain. The License Management Report counts logical CPUs instead of physical CPUs for license enforcement. If the

number of logical CPUs exceeds the number of authorized CPUs, then the License Management Report shows that the domain exceeded the CPU limit.

- Repository usage. Shows the number of PowerCenter Repository Services in the domain.
- User information. Shows information about users in the domain.
- Hardware configuration. Shows details about the machines used in the domain.
- Node configuration. Shows details about each node in the domain.
- Licensed options. Shows a list of PowerCenter and other Informatica options purchased for each license.

Licensing

The Licensing section of the License Management Report shows information about each license in the domain.

The following table describes the licensing information in the License Management Report:

Property	Description
Name	Name of the license.
Edition	PowerCenter edition.
Version	Version of Informatica platform.
Expiration Date	Date when the license expires.
Serial Number	Serial number of the license. The serial number identifies the customer or project. If the customer has multiple PowerCenter installations, there is a separate serial number for each project. The original and incremental keys for a license have the same serial number.
Deployment Level	Level of deployment. Values are Development and Production.
Operating System / BitMode	Operating system and bitmode for the license. Indicates whether the license is installed on a 32-bit or 64-bit operating system.
CPU	Maximum number of authorized logical CPUs.
Repository	Maximum number of authorized PowerCenter repositories.
AT Named Users	Maximum number of users who are assigned the License Access for Informatica Analyst privilege.
Product Bitmode	Bitmode of the server binaries that are installed. Values are 32-bit or 64-bit.

CPU Summary

The CPU Summary section of the License Management Report shows the maximum number of logical CPUs used to run application services in the domain. Use the CPU summary information to determine if the CPU usage exceeded the license limits. If the number of logical CPUs is greater than the total number of CPUs authorized by the license, the License Management Report indicates that the CPU limit is exceeded.

The License Management Report determines the number of logical CPUs based on the number of processors, cores, and threads. Use the following formula to calculate the number of logical CPUs:

$N * C * T$, where

N is the number of processors.

C is the number of cores in each processor.

T is the number of threads in each core.

For example, a machine contains 4 processors. Each processor has 2 cores. The machine contains 8 (4*2) physical cores. Hyperthreading is enabled, where each core contains 3 threads. The number of logical CPUs is 24 (4*2*3).

Note: Although the License Management Report includes threads in the calculation of logical CPUs, Informatica license compliance is based on the number of physical cores, not threads. To be compliant, the number of physical cores must be less than or equal to the maximum number of licensed CPUs. If the License Management Report shows that you have exceeded the license limit but the number of physical cores is less than or equal to the maximum number of licensed CPUs, you can ignore the message. If you have a concern about license compliance, contact your Informatica account manager.

The following table describes the CPU summary information in the License Management Report:

Property	Description
Domain	Name of the domain on which the report runs.
Current Usage	Maximum number of logical CPUs used concurrently on the day the report runs.
Peak Usage	Maximum number of logical CPUs used concurrently during the last 12 months.
Peak Usage Date	Date when the maximum number of logical CPUs were used concurrently during the last 12 months.
Days Exceeded License Limit	Number of days that the CPU usage exceeded the license limits. The domain exceeds the CPU license limit when the number of concurrent logical CPUs exceeds the number of authorized CPUs.

CPU Detail

The CPU Detail section of the License Management Report provides CPU usage information for each host in the domain. The CPU Detail section shows the maximum number of logical CPUs used each day in a selected time period.

The report counts the number of logical CPUs on each host that runs application services in the domain. The report groups logical CPU totals by node.

The following table describes the CPU detail information in the License Management Report:

Property	Description
Host Name	Host name of the machine.
Current Usage	Maximum number of logical CPUs that the host used concurrently on the day the report runs.
Peak Usage	Maximum number of logical CPUs that the host used concurrently during the last 12 months.
Peak Usage Date	Date in the last 12 months when the host concurrently used the maximum number of logical CPUs.
Assigned Licenses	Name of all licenses assigned to services that run on the node.

Repository Summary

The Repository Summary section of the License Management Report provides repository usage information for the domain. Use the repository summary information to determine if the repository usage exceeded the license limits.

The following table describes the repository summary information in the License Management Report:

Property	Description
Current Usage	Maximum number of repositories used concurrently in the domain on the day the report runs.
Peak Usage	Maximum number of repositories used concurrently in the domain during the last 12 months.
Peak Usage Date	Date in the last 12 months when the maximum number of repositories were used concurrently.
Days Exceeded License Limit	Number of days that the repository usage exceeded the license limits.

User Summary

The User Summary section of the License Management Report provides information about Analyst tool users in the domain.

The following table describes the user summary information in the License Management Report:

Property	Description
User Type	Type of user in the domain.
Current Named Users	Maximum number of users who are assigned the License Access for Informatica Analyst privilege on the day the report runs.
Peaked Name Users	Maximum number of users who are assigned the License Access for Informatica Analyst privilege during the last 12 months.
Peak Named Users Date	Date during the last 12 months when the maximum number of concurrent users were assigned the License Access for Informatica Analyst privilege.

User Detail

The User Detail section of the License Management Report provides information about each Analyst tool user in the domain.

The following table describes the user detail information in the License Management Report:

Property	Description
User Type	Type of user in the domain.
User Name	User name.

Property	Description
Days Logged In	Number of days the user logged in to the Analyst tool and performed profiling during the last 12 months.
Peak Unique IP Addresses in a Day	Maximum number of machines that the user was logged in to and performed profiling on during a single day of the last 12 months.
Average Unique IP Addresses	Daily average number of machines that the user was logged in to and running profiling on during the last 12 months.
Peak IP Address Date	Date when the user logged in to and performed profiling on the maximum number of machines during a single day of the last 12 months.
Peak Daily Sessions	Maximum number of times in a single day of the last 12 months that the user logged in to any Analyst tool and performed profiling.
Average Daily Sessions	Average number of times per day in the last 12 months that the user logged in to any Analyst tool and performed profiling.
Peak Session Date	Date in the last 12 months when the user had the most daily sessions in the Analyst tool.

Hardware Configuration

The Hardware Configuration section of the License Management Report provides details about machines used in the domain.

The following table describes the hardware configuration information in the License Management Report:

Property	Description
Host Name	Host name of the machine.
Logical CPUs	Number of logical CPUs used to run application services in the domain.
Sockets	Number of sockets on the machine.
Consumed cores	Number of cores on the machine.
Cores per socket	Number of cores for each socket on the machine.
CPU Model	Model of the CPU.
Hyperthreading Enabled	Indicates whether hyperthreading is enabled.
Virtual Machine	Indicates whether the machine is a virtual machine.

Node Configuration

The Node Configuration section of the License Management Report provides details about each node in the domain.

The following table describes the node configuration information in the License Management Report:

Property	Description
Node Name	Name of the node or nodes assigned to a machine for a license.
Host Name	Host name of the machine.
IP Address	IP address of the node.
Operating System	Operating system of the machine on which the node runs.
Status	Status of the node.
Gateway	Indicates whether the node is a gateway node.
Service Type	Type of the application service configured to run on the node.
Service Name	Name of the application service configured to run on the node.
Service Status	Status of the application service.
Assigned License	License assigned to the application service.

Licensed Options

The Licensed Options section of the License Management Report provides details about each option for every license assigned to the domain.

The following table describes the licensed option information in the License Management Report:

Property	Description
License Name	Name of the license.
Description	Name of the license option.
Status	Status of the license option.
Issued On	Date when the license option was issued.
Expires On	Date when the license option expires.

Running the License Management Report

Run the License Management Report from the **Reports** tab in the Administrator tool.

1. Click the **Reports** tab in the Administrator tool.
2. Click the **License Management Report** view.

The License Management Report appears.

3. Click **Save** to save the License Management Report as a PDF.

If a License Management Report contains multibyte characters, you must configure the Service Manager to use a Unicode font.

- Click **Email** to send a copy of the License Management Report in an email.
The **Send License Management Report** page appears.

Configuring a Unicode Font for the Report

Before you can save a License Management Report that contains multibyte characters or non-English characters, configure the Service Manager to use a Unicode font when generating the PDF file.

- Install a Unicode font on the master gateway node.
- Use a text editor to create a file named `AcUtil.properties`.
- Add the following properties to the file:

```
PDF.Font.Default=Unicode_font_name
PDF.Font.MultibyteList=Unicode_font_name
```

Unicode_font_name is the name of the Unicode font installed on the master gateway node.

You might also need to add the following property if the font file is not available in the locale:

```
Unicode_font_name_path=Unicode_font_file_location
```

For example:

```
PDF.Font.Default=Arial Unicode MS
PDF.Font.MultibyteList=Arial Unicode MS
Arial Unicode MS_path=/usr/lib/X11/fonts/TrueType
```

- Save the `AcUtil.properties` file to the following location:

```
InformaticaInstallationDir\services\AdministratorConsole\administrator
```

- Use a text editor to open the `licenseUtility.css` file in the following location:

```
InformaticaInstallationDir\services\AdministratorConsole\administrator\css
```

- Append the Unicode font name to the value of each font-family property.

For example:

```
font-family: Arial Unicode MS, Verdana, Arial, Helvetica, sans-serif;
```

- Restart Informatica services on each node in the domain.

Sending the License Management Report in an Email

You must configure the SMTP settings for the domain before you can send the License Management Report in an email.

The domain administrator can send the License Management Report in an email from Send License Management Report page in the Administrator tool.

- Enter the following information:

Property	Description
To Email	Email address to which you send the License Management Report.
Subject	Subject of the email.
Customer Name	Name of the organization that purchased the license.

Property	Description
Request ID	Request ID that identifies the project for which the license was purchased.
Contact Name	Name of the contact person in the organization.
Contact Phone Number	Phone number of the contact person.
Contact Email	Email address of the contact person at the customer site.

2. Click OK.

The Administrator tool sends the License Management Report in an email.

Web Services Report

To analyze the performance of web services running on a Web Services Hub, you can run a report for the Web Services Hub or for a web service running on the Web Services Hub.

The Web Services Report provides run-time and historical information on the web service requests handled by the Web Services Hub. The report displays aggregated information for all web services in the Web Services Hub and information for each web service running on the Web Services Hub. The Web Services Report also provides historical information.

Understanding the Web Services Report

You can run the Web Services Report for a time interval that you choose. The Web Services Hub collects information on web services activities and caches 24 hours of information for use in the Web Services Report. It also writes the information to a history file.

Time Interval

By default, the Web Services Report displays activity information for a five-minute interval. You can select one of the following time intervals to display activity information for a web service or Web Services Hub:

- 5 seconds
- 1 minute
- 5 minutes
- 1 hour
- 24 hours

The Web Services Report displays activity information for the interval ending at the time you run the report. For example, if you run the Web Services Report at 8:05 a.m. for an interval of one hour, the Web Services Report displays the Web Services Hub activity from 7:05 a.m. and 8:05 a.m.

Caching

The Web Services Hub caches 24 hours of activity data. The cache is reinitialized every time the Web Services Hub is restarted. The Web Services Report displays statistics from the cache for the time interval that you run the report.

History File

The Web Services Hub writes the cached activity data to a history file. The Web Services Hub stores data in the history file for the number of days that you set in the MaxStatsHistory property of the Web Services Hub. For example, if the value of the MaxStatsHistory property is 5, the Web Services Hub keeps five days of data in the history file.

Contents of the Web Services Report

The Web Services Report view contains information about the web services in the domain. When you select a web services hub in the Navigator, you can view the following information about the web services it contains:

- Properties view. Displays General Properties, Web Services Hub Summary, and Historical Statistics for the web services hub.
- Web Services view. Lists the web services in the web services hub. When you select a web service, you can view Properties, Top IP Addresses, and Historical Statistics for the web service.

General Properties and Web Services Hub Summary

To view the general properties and summary information for the Web Services Hub, select the Properties view in the content panel.

The following table describes the general properties:

Property	Description
Name	Name of the Web Services Hub.
Description	Short description of the Web Services Hub.
Service type	Type of Service. For a Web Services Hub, the service type is ServiceWSHubService.

The following table describes the Web Services Hub Summary properties:

Property	Description
# of Successful Message	Number of requests that the Web Services Hub processed successfully.
# of Fault Responses	Number of fault responses generated by web services in the Web Services Hub. The fault responses could be due to any error.
Total Messages	Total number of requests that the Web Services Hub received.
Last Server Restart Tme	Date and time when the Web Services Hub was last started.
Avg. # of Service Partitions	Average number of partitions allocated for all web services in the Web Services Hub.
% of Partitions in Use	Percentage of web service partitions that are in use for all web services in the Web Services Hub.
Avg. # of Run Instances	Average number of instances running for all web services in the Web Services Hub.

Web Services Historical Statistics

To view historical statistics for the web services in the Web Services Hub, select the Properties view in the content panel. The detail panel displays data from the Web Services Hub history file for the date that you specify.

The following table describes the historical statistics:

Property	Description
Time	Time of the event.
Web Service	Name of the web service for which the information is displayed. When you click the name of a web service, the Web Services Report displays the Service Statistics window.
Successful Requests	Number of requests successfully processed by the web service.
Fault Responses	Number of fault responses sent by the web service.
Avg. Service Time	Average time it takes to process a service request received by the web service.
Max Service Time	The largest amount of time taken by the web service to process a request.
Min Service Time	The smallest amount of time taken by the web service to process a request.
Avg. DTM Time	Average number of seconds it takes the PowerCenter Integration Service to process the requests from the Web Services Hub.
Avg. Service Partitions	Average number of session partitions allocated for the web service.
Percent Partitions in Use	Percentage of partitions in use by the web service.
Avg Run Instances	Average number of instances running for the web service.

Web Services Run-time Statistics

To view run-time statistics for each web service in the Web Services Hub, select the Web Services view in the content panel. The Web Services view lists the statistics for each web service.

The report provides the following information for each web service for the selected time interval:

Property	Description
Service name	Name of the web service for which the information is displayed.
Successful Requests	Number of requests received by the web service that the Web Services Hub processed successfully.
Fault Responses	Number of fault responses generated by the web services in the Web Services Hub.
Avg. Service Time	Average time it takes to process a service request received by the web service.

Property	Description
Avg. Service Partitions	Average number of session partitions allocated for the web service.
Avg. Run Instances	Average number of instances of the web service running during the interval.

Web Service Properties

To view the properties of a web service, select the web service in the Web Services view of the content panel. In the details panel, the Properties view displays the properties for the web service.

The report provides the following information for the selected web service:

Property	Description
# of Successful Requests	Number of requests received by the web service that the Web Services Hub processed successfully.
# of Fault Responses	Number of fault responses generated by the web services in the Web Services Hub.
Total Messages	Total number of requests that the Web Services Hub received.
Last Server Restart Time	Date and time when the Web Services Hub was last started
Last Service Time	Number of seconds it took to process the most recent service request
Average Service Time	Average time it takes to process a service request received by the web service.
Avg. # of Service Partitions	Average number of session partitions allocated for the web service.
Avg. # of Run Instances	Average number of instances of the web service running during the interval.

Web Service Top IP Addresses

To view the top IP addresses for a web service, select a web service in the Web Services view of the content panel and select the Top IP Addresses view in the details panel. The Top IP Addresses displays the most active IP addresses for the web service, listed in the order of longest to shortest service times.

The report provides the following information for each of the most active IP addresses:

Property	Description
Top 10 Client IP Addresses	The list of client IP addresses and the longest time taken by the web service to process a request from the client. The client IP addresses are listed in the order of longest to shortest service times. Use the Click here link to display the list of IP addresses and service times.

Web Service Historical Statistics Table

To view a table of historical statistics for a web service, select a web service in the Web Services view of the content panel and select the Table view in the details panel. The details panel displays a table of historical statistics for the web service.

The table provides the following information for the selected web service:

Property	Description
Time	Time of the event.
Web Service	Name of the web service for which the information is displayed.
Successful Requests	Number of requests successfully processed by the web service.
Fault Responses	Number of requests received for the web service that could not be processed and generated fault responses.
Avg. Service Time	Average time it takes to process a service request received by the web service.
Min. Service Time	The smallest amount of time taken by the web service to process a request.
Max. Service Time	The largest amount of time taken by the web service to process a request.
Avg. DTM Time	Average time it takes the PowerCenter Integration Service to process the requests from the Web Services Hub.
Avg. Service Partitions	Average number of session partitions allocated for the web service.
Percent Partitions in Use	Percentage of partitions in use by the web service.
Avg. Run Instances	Average number of instances running for the web service.

Running the Web Services Report

Run the Web Services Report from the Reports tab in the Administrator tool.

Before you run the Web Services Report for a Web Services Hub, verify that the Web Services Hub is enabled. You cannot run the Web Services Report for a disabled Web Services Hub.

1. In the Administrator tool, click the Reports tab.
2. Click Web Services.
3. In the Navigator, select the Web Services Hub for which to run the report.
In the content panel, the Properties view displays the properties of the Web Services Hub. The details view displays historical statistics for the services in the Web Services Hub.
4. To specify a date for historical statistics, click the date filter icon in the details panel, and select the date.
5. To view information about each service, select the Web Services view in the content panel.
The Web Services view displays summary statistics for each service for the Web Services Hub.
6. To view additional information about a service, select the service from the list.
In the details panel, the Properties view displays the properties for the service.
7. To view top IP addresses for the service, select the Top IP Addresses view in the details panel.
8. To view table attributes for the service, select the Table view in the detail panel.

Running the Web Services Report for a Secure Web Services Hub

To run a Web Services Hub on HTTPS, you must have an SSL certificate file for authentication of message transfers. When you create a Web Services Hub to run on HTTPS, you must specify the location of the keystore file that contains the certificate for the Web Services Hub. To run the Web Services Report in the Administrator tool for a secure Web Services Hub, you must import the SSL certificate into the Java certificate file. The Java certificate file is named *cacerts* and is located in the */lib/security* directory of the Java directory. The Administrator tool uses the *cacerts* certificate file to determine whether to trust an SSL certificate.

In a domain that contains multiple nodes, the node where you generate the SSL certificate affects how you access the Web Services Report for a secure Web Services Hub.

Use the following rules and guidelines to run the Web Services Report for a secure Web Services Hub in a domain with multiple nodes:

- For each secure Web Services Hub running in a domain, generate an SSL certificate and import it to a Java certificate file.
- The Administrator tool searches for SSL certificates in the certificate file of a gateway node. The SSL certificate for a Web Services Hub running on worker node must be generated on a gateway node and imported into the certificate file of the same gateway node.
- To view the Web Services Report for a secure Web Services Hub, log in to the Administrator tool from the gateway node that has the certificate file containing the SSL certificate of the Web Services Hub for which you want to view reports.
- If a secure Web Services Hub runs on a worker node, the SSL certificate must be generated and imported into the certificate file of the gateway node. If a secure Web Services Hub runs on a gateway and a worker node, the SSL certificate of both nodes must be generated and imported into the certificate file of the gateway node. To view reports for the secure Web Services Hub, log in to the Administrator tool from the gateway node.
- If the domain has two gateway nodes and a secure Web Services Hub runs on each gateway node, access to the Web Services Reports depends on where the SSL certificate is located.

For example, gateway node GWN01 runs Web Services Hub WSH01 and gateway node GWN02 runs Web Services Hub WSH02. You can view the reports for the Web Services Hubs based on the location of the SSL certificates:

- If the SSL certificate for WSH01 is in the certificate file of GWN01 but not GWN02, you can view the reports for WSH01 if you log in to the Administrator tool through GWN01. You cannot view the reports for WSH01 if you log in to the Administrator tool through GWN02. If GWN01 fails, you cannot view reports for WSH01.
- If the SSL certificate for WSH01 is in the certificate files of GWN01 and GWN02, you can view the reports for WSH01 if you log in to the Administrator tool through GWN01 or GWN02. If GWN01 fails, you can view the reports for WSH01 if you log in to the Administrator tool through GWN02.
- To ensure successful failover when a gateway node fails, generate and import the SSL certificates of all Web Services Hubs in the domain into the certificates files of all gateway nodes in the domain.

CHAPTER 16

Node Diagnostics

This chapter includes the following topics:

- [Node Diagnostics Overview, 294](#)
- [Informatica Network Login, 295](#)
- [Generating Node Diagnostics, 296](#)
- [Downloading Node Diagnostics, 297](#)
- [Uploading Node Diagnostics, 297](#)
- [Analyzing Node Diagnostics, 298](#)

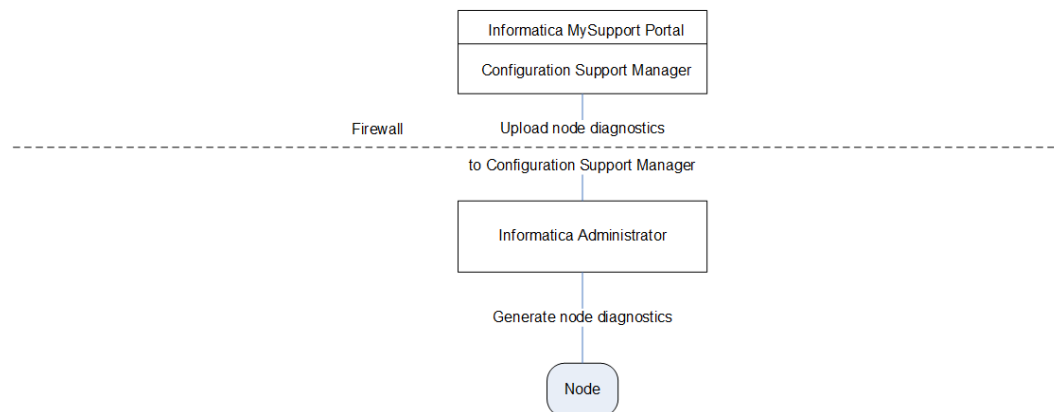
Node Diagnostics Overview

The Configuration Support Manager is a web-based application that you can use to track Informatica updates and diagnose issues in your environment.

You can discover comprehensive information about your technical environment and diagnose issues before they become critical.

Generate node diagnostics from the Informatica Administrator and upload them to the Configuration Support Manager in the Informatica MySupport Portal. Then, check the node diagnostics against business rules and recommendations in the Configuration Support Manager.

The following image shows the operational flow to generate and upload node diagnostics:



Complete the following tasks to generate and upload node diagnostics:

1. Log in to the Informatica MySupport Portal.

2. Generate node diagnostics. The Service Manager analyzes the services of the node and generates node diagnostics including information such as operating system details, CPU details, database details, and patches.
3. Optionally, download node diagnostics to your local drive.
4. Upload node diagnostics to the Configuration Support Manager, a diagnostic web application outside the firewall. The Configuration Support Manager is a part of the Informatica MySupport Portal. The Service Manager connects to the Configuration Support Manager through the HTTPS protocol and uploads the node diagnostics.
5. Review the node diagnostics in the Configuration Support Manager to find troubleshooting information for your environment.

Informatica Network Login

You must log in to the Informatica Network to upload node diagnostics to the Configuration Support Manager. The login credentials are not specific to a user. The same credentials are applicable for all users who have access to the Administrator tool. Register at <http://communities.informatica.com> if you do not have the customer portal login details. You need to enter the customer portal login details and save these details. Alternatively, you can enter the customer portal details each time you upload node diagnostics to the Configuration Support Manager. You can generate node diagnostics without entering the login details.

To maintain login security, you must log out of the Configuration Support Manager and the Node Diagnostics Upload page of the Administrator tool.

- To log out of the Configuration Support Manager, click the logout link.
- To log out of the Upload page, click **Close Window**.

Note: If you close these windows through the web browser close button, you remain logged in to the Configuration Support Manager. Other users can access the Configuration Support Manager without valid credentials.


Logging In to the Informatica Network

Before you generate and upload node diagnostics, you must log in to the Informatica Network.

1. In the Administrator tool, click the **Manage** tab > **Services and Nodes** view.
2. In the Domain Navigator, select the domain.
3. In the contents panel, click **Diagnostics**.

A list of all the nodes in the domain appears.

- Click the **Actions** menu in the upper right corner of the page and select **Edit Informatica Network Login Credentials**:



The **Edit Informatica Network Login Credentials** dialog box appears.

- Enter the following customer portal login details:

Field	Description
Email Address	Email address with which you registered your customer portal account.
Password	Password for your customer portal account.
Project ID	Unique ID assigned to your support project.

- Click **OK**.

Generating Node Diagnostics

When you generate node diagnostics, the Administrator tool generates node diagnostics in an XML file.

The XML file contains details about services, logs, environment variables, operating system parameters, system information, and database clients. Worker node diagnostics only contain node metadata; they do not include domain metadata.

- In the Administrator tool, click the **Manage** tab > **Services and Nodes** view.
- In the Domain Navigator, select the domain.
- In the contents panel, click **Diagnostics**.
A list of all nodes in the domain appears.
- Select the node.
- Click **Generate Diagnostics File**.
- Click **Yes** to confirm that you want to generate node diagnostics.

Note: You can also generate diagnostics from the **Actions** menu on the **Diagnostics** tab.

The csmagent<host name>.xml file, which contains the node diagnostics, is generated at `INFA_HOME/server/csm/output`. The node diagnostics and the time stamp of the generated file appear.

7. To run diagnostics for your environment, upload the csmagent<host name>.xml file to the Configuration Support Manager.

Alternatively, you can download the XML file to your local drive.

After you generate node diagnostics for the first time, you can regenerate or upload them.

Downloading Node Diagnostics

After you generate node diagnostics, you can download the file to your local drive.

1. In the Administrator tool, click the **Manage** tab > **Services and Nodes** view.
2. In the Domain Navigator, select the domain.
3. In the contents panel, click **Diagnostics**.
A list of all nodes in the domain appears.
4. Click the diagnostics file name of the node.
The file opens in another browser window.
5. Click **File** > **Save As**. Then, specify a location to save the file.
6. Click **Save**.
The XML file is saved to your local drive.

Uploading Node Diagnostics

You can upload node diagnostics to the Configuration Support Manager through the Administrator tool. You must enter the customer portal login details before you upload node diagnostics.

When you upload node diagnostics, you can update or create a configuration in the Configuration Support Manager. Create a configuration the first time you upload the node diagnostics. Update a configuration to view the latest diagnostics of the configuration. To compare current and previous node configurations of an existing configuration, upload the current node diagnostics as a new configuration.

Note: If you do not have access to the Internet, you can download the file and upload it at a later time. You can also send the file to the Informatica Global Customer Support in an email to troubleshoot or to upload.

1. In the Administrator tool, click the **Manage** tab > **Services and Nodes** view.
2. In the Domain Navigator, select the domain.
3. In the contents panel, click **Diagnostics**.
A list of all nodes in the domain appears.
4. Select the node.
5. Generate node diagnostics.
6. Click **Upload Diagnostics File to CSM**.

You can upload the node diagnostics as a new configuration or as an update to an existing configuration.

- To upload a new configuration, go to step [10](#).
To update a configuration, select **Update an existing configuration**.
- Select the configuration you want to update from the list of configurations.
- Go to step [12](#).
- Select **Upload as a new configuration**.
- Enter the following configuration details:

Field	Description
Name	Configuration name.
Description	Configuration description.
Type	Type of the node, which is one of the following types: <ul style="list-style-type: none">- Production- Development- Test/QA

- Click **Upload Now**.

After you upload the node diagnostics, go to the Configuration Support Manager to analyze the node diagnostics.

- Click **Close Window**.

Note: If you close the window by using the close button in the browser, the user authentication session does not end and you cannot upload node diagnostics to the Configuration Support Manager with another set of customer portal login credentials.

Analyzing Node Diagnostics

Use the Configuration Support Manager to analyze node diagnostics.

Use the Configuration Support Manager to complete the following tasks:

- Diagnose issues before they become critical.
- Identify bug fixes.
- Identify recommendations that can reduce risk of unplanned outage.
- View details of your technical environment.
- Manage your configurations efficiently.
- Subscribe to proactive alerts through email and RSS.
- Run advanced diagnostics with compare configuration.

Identify Bug Fixes

You can use the Configuration Support Manager to resolve issues encountered during operations. To expedite resolution of support issues, you can generate and upload node diagnostics to the Configuration

Support Manager. You can analyze node diagnostics in the Configuration Support Manager and find a solution to your issue.

For example, when you run a Sorter session that processes a large volume of data, you notice that there is some data loss. You generate node diagnostics and upload them to the Configuration Support Manager. When you review the diagnostics for bug fix alerts, you see that a bug fix, EBF178626, is available for this. You apply EBF178626, and run the session again. All data is successfully loaded.

Identify Recommendations

You can use the Configuration Support Manager to avoid issues in your environment. You can troubleshoot issues that arise after you make changes to the node properties by comparing different node diagnostics in the Configuration Support Manager. You can also use the Configuration Support Manager to identify recommendations or updates that may help you improve the performance of the node.

For example, you upgrade the node memory to handle a higher volume of data. You generate node diagnostics and upload them to the Configuration Support Manager. When you review the diagnostics for operating system warnings, you find the recommendation to increase the total swap memory of the node to twice that of the node memory for optimal performance. You increase swap space as suggested in the Configuration Support Manager and avoid performance degradation.

Tip: Regularly upload node diagnostics to the Configuration Support Manager and review node diagnostics to maintain your environment efficiently.

CHAPTER 17

Understanding Globalization

This chapter includes the following topics:

- [Globalization Overview, 300](#)
- [Locales, 302](#)
- [Data Movement Modes, 303](#)
- [Code Page Overview, 305](#)
- [Code Page Compatibility, 307](#)
- [Code Page Validation, 314](#)
- [Relaxed Code Page Validation, 315](#)
- [PowerCenter Code Page Conversion, 317](#)
- [Case Study: Processing ISO 8859-1 Data, 318](#)
- [Case Study: Processing Unicode UTF-8 Data, 320](#)

Globalization Overview

Informatica can process data in different languages. Some languages require single-byte data, while other languages require multibyte data. To process data correctly in Informatica, you must set up the following items:

- **Locale.** Informatica requires that the locale settings on machines that access Informatica applications are compatible with code pages in the domain. You may need to change the locale settings. The locale specifies the language, territory, encoding of character set, and collation order.
- **Data movement mode.** The PowerCenter Integration Service can process single-byte or multibyte data and write it to targets. Use the ASCII data movement mode to process single-byte data. Use the Unicode data movement mode for multibyte data.
- **Code pages.** Code pages contain the encoding to specify characters in a set of one or more languages. You select a code page based on the type of character data you want to process. To ensure accurate data movement, you must ensure compatibility among code pages for Informatica and environment components. You use code pages to distinguish between US-ASCII (7-bit ASCII), ISO 8859-1 (8-bit ASCII), and multibyte characters.

To ensure data passes accurately through your environment, the following components must work together:

- Domain configuration database code page
- Administrator tool locale settings and code page
- PowerCenter Integration Service data movement mode

- Code page for each PowerCenter Integration Service process
- PowerCenter Client code page
- PowerCenter repository code page
- Source and target database code pages
- Metadata Manager repository code page

You can configure the PowerCenter Integration Service for relaxed code page validation. Relaxed validation removes restrictions on source and target code pages.

Unicode

The Unicode Standard is the work of the Unicode Consortium, an international body that promotes the interchange of data in all languages. The Unicode Standard is designed to support any language, no matter how many bytes each character in that language may require. Currently, it supports all common languages and provides limited support for other less common languages. The Unicode Consortium is continually enhancing the Unicode Standard with new character encodings. For more information about the Unicode Standard, see <http://www.unicode.org>.

The Unicode Standard includes multiple character sets. Informatica uses the following Unicode standards:

- UCS-2 (Universal Character Set, double-byte). A character set in which each character uses two bytes.
- UTF-8 (Unicode Transformation Format). An encoding format in which each character can use between one to four bytes.
- UTF-16 (Unicode Transformation Format). An encoding format in which each character uses two or four bytes.
- UTF-32 (Unicode Transformation Format). An encoding format in which each character uses four bytes.
- GB18030. A Unicode encoding format defined by the Chinese government in which each character can use between one to four bytes.

Informatica is a Unicode application. The PowerCenter Client, PowerCenter Integration Service, and Data Integration Service use UCS-2 internally. The PowerCenter Client converts user input from any language to UCS-2 and converts it from UCS-2 before writing to the PowerCenter repository. The PowerCenter Integration Service and Data Integration Service converts source data to UCS-2 before processing and converts it from UCS-2 after processing. The PowerCenter repository, Model repository, PowerCenter Integration Service, and Data Integration Service support UTF-8. You can use Informatica to process data in any language.

Working with a Unicode PowerCenter Repository

The PowerCenter repository code page is the code page of the data in the PowerCenter repository. You choose the PowerCenter repository code page when you create or upgrade a PowerCenter repository. When the PowerCenter repository database code page is UTF-8, you can create a PowerCenter repository using the UTF-8 code page.

The domain configuration database uses the UTF-8 code page. If you need to store metadata in multiple languages, such as Chinese, Japanese, and Arabic, you must use the UTF-8 code page for all services in that domain.

The Service Manager synchronizes the list of users in the domain with the list of users and groups in each application service. If a user in the domain has characters that the code page of the application services does not recognize, characters do not convert correctly and inconsistencies occur.

Use the following guidelines when you use UTF-8 as the PowerCenter repository code page:

- The PowerCenter repository database code page must be UTF-8.

- The PowerCenter repository code page must be a superset of the PowerCenter Client and PowerCenter Integration Service process code pages.
- You can input any character in the UCS-2 character set. For example, you can store German, Chinese, and English metadata in a UTF-8 enabled PowerCenter repository.
- Install languages and fonts on the PowerCenter Client machine. If you are using a UTF-8 PowerCenter repository, you may want to enable the PowerCenter Client machines to display multiple languages. By default, the PowerCenter Clients display text in the language set in the system locale. Use the Regional Options tool in the Control Panel to add language groups to the PowerCenter Client machines.
- You can use the Windows Input Method Editor (IME) to enter multibyte characters from any language without having to run the version of Windows specific for that language.
- Choose a code page for a PowerCenter Integration Service process that can process all PowerCenter repository metadata correctly. The code page of the PowerCenter Integration Service process must be a subset of the PowerCenter repository code page. If the PowerCenter Integration Service has multiple service processes, ensure that the code pages for all PowerCenter Integration Service processes are subsets of the PowerCenter repository code page. If you are running the PowerCenter Integration Service process on Windows, the code page for the PowerCenter Integration Service process must be the same as the code page for the system or user locale. If you are running the PowerCenter Integration Service process on UNIX, use the UTF-8 code page for the PowerCenter Integration Service process.

Locales

Every machine has a locale. A locale is a set of preferences related to the user environment, including the input language, keyboard layout, how data is sorted, and the format for currency and dates. Informatica uses locale settings on each machine.

You can set the following locale settings on Windows:

- System locale. Determines the language, code pages, and associated bitmap font files that are used as defaults for the system.
- User locale. Determines the default formats to display date, time, currency, and number formats.
- Input locale. Describes the input method, such as the keyboard, of the system language.

For more information about configuring the locale settings on Windows, consult the Windows documentation.

System Locale

The system locale is also referred to as the system default locale. It determines which ANSI and OEM code pages, as well as bitmap font files, are used as defaults for the system. The system locale contains the language setting, which determines the language in which text appears in the user interface, including in dialog boxes and error messages. A message catalog file defines the language in which messages display. By default, the machine uses the language specified for the system locale for all processes, unless you override the language for a specific process.

The system locale is already set on your system and you may not need to change settings to run Informatica. If you do need to configure the system locale, you configure the locale on a Windows machine in the Regional Options dialog box. On UNIX, you specify the locale in the LANG environment variable.

User Locale

The user locale displays date, time, currency, and number formats for each user. You can specify different user locales on a single machine. Create a user locale if you are working with data on a machine that is in a different language than the operating system. For example, you might be an English user working in Hong Kong on a Chinese operating system. You can set English as the user locale to use English standards in your work in Hong Kong. When you create a new user account, the machine uses a default user locale. You can change this default setting once the account is created.

Input Locale

An input locale specifies the keyboard layout of a particular language. You can set an input locale on a Windows machine to type characters of a specific language.

You can use the Windows Input Method Editor (IME) to enter multibyte characters from any language without having to run the version of Windows specific for that language. For example, if you are working on an English operating system and need to enter text in Chinese, you can use IME to set the input locale to Chinese without having to install the Chinese version of Windows. You might want to use an input method editor to enter multibyte characters into a PowerCenter repository that uses UTF-8.

Data Movement Modes

The data movement mode is a PowerCenter Integration Service option you choose based on the type of data you want to move, single-byte or multibyte data. The data movement mode you select depends the following factors:

- Requirements to store single-byte or multibyte metadata in the PowerCenter repository
- Requirements to access source data containing single-byte or multibyte character data
- Future needs for single-byte and multibyte data

The data movement mode affects how the PowerCenter Integration Service enforces session code page relationships and code page validation. It can also affect performance. Applications can process single-byte characters faster than multibyte characters.

Character Data Movement Modes

The PowerCenter Integration Service runs in the following modes:

- ASCII (American Standard Code for Information Interchange). The US-ASCII code page contains a set of 7-bit ASCII characters and is a subset of other character sets. When the PowerCenter Integration Service runs in ASCII data movement mode, each character requires one byte.
- Unicode. The universal character-encoding standard that supports all languages. When the PowerCenter Integration Service runs in Unicode data movement mode, it allots up to two bytes for each character. Run the PowerCenter Integration Service in Unicode mode when the source contains multibyte data.

Tip: You can also use ASCII or Unicode data movement mode if the source has 8-bit ASCII data. The PowerCenter Integration Service allots an extra byte when processing data in Unicode data movement mode. To increase performance, use the ASCII data movement mode. For example, if the source contains characters from the ISO 8859-1 code page, use the ASCII data movement.

The data movement you choose affects the requirements for code pages. Ensure the code pages are compatible.

ASCII Data Movement Mode

In ASCII mode, the PowerCenter Integration Service processes single-byte characters and does not perform code page conversions. When you run the PowerCenter Integration Service in ASCII mode, it does not enforce session code page relationships.

Unicode Data Movement Mode

In Unicode mode, the PowerCenter Integration Service recognizes multibyte character data and allocates up to two bytes for every character. The PowerCenter Integration Service performs code page conversions from sources to targets. When you set the PowerCenter Integration Service to Unicode data movement mode, it uses a Unicode character set to process characters in a specified code page, such as Shift-JIS or UTF-8.

When you run the PowerCenter Integration Service in Unicode mode, it enforces session code page relationships.

Changing Data Movement Modes

You can change the data movement mode in the PowerCenter Integration Service properties in the Administrator tool. After you change the data movement mode, the PowerCenter Integration Service runs in the new data movement mode the next time you start the PowerCenter Integration Service. When the data movement mode changes, the PowerCenter Integration Service handles character data differently. To avoid creating data inconsistencies in your target tables, the PowerCenter Integration Service performs additional checks for sessions that reuse session caches and files.

The following table describes how the PowerCenter Integration Service handles session files and caches after you change the data movement mode:

Session File or Cache	Time of Creation or Use	PowerCenter Integration Service Behavior After Data Movement Mode Change
Session Log File (*.log)	Each session.	No change in behavior. Creates a new session log for each session using the code page of the PowerCenter Integration Service process.
Workflow Log	Each workflow.	No change in behavior. Creates a new workflow log file for each workflow using the code page of the PowerCenter Integration Service process.
Reject File (*.bad)	Each session.	No change in behavior. Appends rejected data to the existing reject file using the code page of the PowerCenter Integration Service process.
Output File (*.out)	Sessions writing to flat file.	No change in behavior. Creates a new output file for each session using the target code page.
Indicator File (*.in)	Sessions writing to flat file.	No change in behavior. Creates a new indicator file for each session.

Session File or Cache	Time of Creation or Use	PowerCenter Integration Service Behavior After Data Movement Mode Change
Incremental Aggregation Files (*.idx, *.dat)	Sessions with Incremental Aggregation enabled.	<p>When files are removed or deleted, the PowerCenter Integration Service creates new files.</p> <p>When files are not moved or deleted, the PowerCenter Integration Service fails the session with the following error message:</p> <pre>SM_7038 Aggregate Error: ServerMode: [server data movement mode] and CachedMode: [data movement mode that created the files] mismatch.</pre> <p>Move or delete files created using a different code page.</p>
Unnamed Persistent Lookup Files (*.idx, *.dat)	Sessions with a Lookup transformation configured for an unnamed persistent lookup cache.	Rebuilds the persistent lookup cache.
Named Persistent Lookup Files (*.idx, *.dat)	Sessions with a Lookup transformation configured for a named persistent lookup cache.	<p>When files are removed or deleted, the PowerCenter Integration Service creates new files.</p> <p>When files are not moved or deleted, the PowerCenter Integration Service fails the session.</p> <p>Move or delete files created using a different code page.</p>

Code Page Overview

A code page contains the encoding to specify characters in a set of one or more languages. An encoding is the assignment of a number to a character in the character set. You use code pages to identify data that might be in different languages. For example, if you create a mapping to process Japanese data, you must select a Japanese code page for the source data.

When you choose a code page, the program or application for which you set the code page refers to a specific set of data that describes the characters the application recognizes. This influences the way that application stores, receives, and sends character data.

Most machines use one of the following code pages:

- US-ASCII (7-bit ASCII)
- MS Latin1 (MS 1252) for Windows operating systems
- Latin1 (ISO 8859-1) for UNIX operating systems
- IBM EBCDIC US English (IBM037) for mainframe systems

The US-ASCII code page contains all 7-bit ASCII characters and is the most basic of all code pages with support for United States English. The US-ASCII code page is not compatible with any other code page. When you install either the PowerCenter Client, PowerCenter Integration Service, or PowerCenter repository on a US-ASCII system, you must install all components on US-ASCII systems and run the PowerCenter Integration Service in ASCII mode.

MS Latin1 and Latin1 both support English and most Western European languages and are compatible with each other. When you install the PowerCenter Client, PowerCenter Integration Service, or PowerCenter

repository on a system using one of these code pages, you can install the rest of the components on any machine using the MS Latin1 or Latin1 code pages.

You can use the IBM EBCDIC code page for the PowerCenter Integration Service process when you install it on a mainframe system. You cannot install the PowerCenter Client or PowerCenter repository on mainframe systems, so you cannot use the IBM EBCDIC code page for PowerCenter Client or PowerCenter repository installations.

UNIX Code Pages

In the United States, most UNIX operating systems have more than one code page installed and use the ASCII code page by default. If you want to run PowerCenter in an ASCII-only environment, you can use the ASCII code page and run the PowerCenter Integration Service in ASCII mode.

UNIX systems allow you to change the code page by changing the LANG, LC_CTYPE or LC_ALL environment variable. For example, you want to change the code page an AIX machine uses. Use the following command in the C shell to view your environment:

```
locale
```

This results in the following output, in which "C" implies "ASCII":

```
LANG="C"  
LC_CTYPE="C"  
LC_NUMERIC="C"  
LC_TIME="C"  
LC_ALL="C"
```

To change the language to English and require the system to use the Latin1 code page, you can use the following command:

```
setenv LANG en_US.iso88591
```

When you check the locale again, it has been changed to use Latin1 (ISO 8859-1):

```
LANG="en_US.iso88591"  
LC_CTYPE="en_US.iso88591"  
LC_NUMERIC="en_US.iso88591"  
LC_TIME="en_US.iso88591"  
LC_ALL="en_US.iso88591"
```

For more information about changing the locale or code page of a UNIX system, see the UNIX documentation.

Windows Code Pages

The Windows operating system is based on Unicode, but does not display the code page used by the operating system in the environment settings. However, you can make an educated guess based on the country in which you purchased the system and the language the system uses.

If you purchase Windows in the United States and use English as an input and display language, your operating system code page is MS Latin1 (MS1252) by default. However, if you install additional display or input languages from the Windows installation CD and use those languages, the operating system might use a different code page.

For more information about the default code page for your Windows system, contact Microsoft.

Choosing a Code Page

Choose code pages based on the character data you use in mappings. Character data can be represented by character modes based on the character size. Character size is the storage space a character requires in the database. Different character sizes can be defined as follows:

- Single-byte. A character represented as a unique number between 0 and 255. One byte is eight bits. ASCII characters are single-byte characters.
- Double-byte. A character two bytes or 16 bits in size represented as a unique number 256 or greater. Many Asian languages, such as Chinese, have double-byte characters.
- Multibyte. A character two or more bytes in size is represented as a unique number 256 or greater. Many Asian languages, such as Chinese, have multibyte characters.

Code Page Compatibility

Compatibility between code pages is essential for accurate data movement when the PowerCenter Integration Service runs in the Unicode data movement mode.

A code page can be compatible with another code page, or it can be a subset or a superset of another:

- Compatible. Two code pages are compatible when the characters encoded in the two code pages are virtually identical. For example, JapanEUC and JIPSE code pages contain identical characters and are compatible with each other. The PowerCenter repository and PowerCenter Integration Service process can each use one of these code pages and can pass data back and forth without data loss.
- Superset. A code page is a superset of another code page when it contains all the characters encoded in the other code page and additional characters not encoded in the other code page. For example, MS Latin1 is a superset of US-ASCII because it contains all characters in the US-ASCII code page.
Note: Informatica considers a code page to be a superset of itself and all other compatible code pages.
- Subset. A code page is a subset of another code page when all characters in the code page are also encoded in the other code page. For example, US-ASCII is a subset of MS Latin1 because all characters in the US-ASCII code page are also encoded in the MS Latin1 code page.

For accurate data movement, the target code page must be a superset of the source code page. If the target code page is not a superset of the source code page, the PowerCenter Integration Service may not process all characters, resulting in incorrect or missing data. For example, Latin1 is a superset of US-ASCII. If you select Latin1 as the source code page and US-ASCII as the target code page, you might lose character data if the source contains characters that are not included in US-ASCII.

When you install or upgrade a PowerCenter Integration Service to run in Unicode mode, you must ensure code page compatibility among the domain configuration database, the Administrator tool, PowerCenter Clients, PowerCenter Integration Service process nodes, the PowerCenter repository, the Metadata Manager repository, and the machines hosting *pmrep* and *pmcmd*. In Unicode mode, the PowerCenter Integration Service enforces code page compatibility between the PowerCenter Client and the PowerCenter repository, and between the PowerCenter Integration Service process and the PowerCenter repository. In addition, when you run the PowerCenter Integration Service in Unicode mode, code pages associated with sessions must have the appropriate relationships:

- For each source in the session, the source code page must be a subset of the target code page. The PowerCenter Integration Service does not require code page compatibility between the source and the PowerCenter Integration Service process or between the PowerCenter Integration Service process and the target.

- If the session contains a Lookup or Stored Procedure transformation, the database or file code page must be a subset of the target that receives data from the Lookup or Stored Procedure transformation and a superset of the source that provides data to the Lookup or Stored Procedure transformation.
- If the session contains an External Procedure or Custom transformation, the procedure must pass data in a code page that is a subset of the target code page for targets that receive data from the External Procedure or Custom transformation.

Informatica uses code pages for the following components:

- Domain configuration database. The domain configuration database must be compatible with the code pages of the PowerCenter repository and Metadata Manager repository.
- Administrator tool. You can enter data in any language in the Administrator tool.
- PowerCenter Client. You can enter metadata in any language in the PowerCenter Client.
- PowerCenter Integration Service process. The PowerCenter Integration Service can move data in ASCII mode and Unicode mode. The default data movement mode is ASCII, which passes 7-bit ASCII or 8-bit ASCII character data. To pass multibyte character data from sources to targets, use the Unicode data movement mode. When you run the PowerCenter Integration Service in Unicode mode, it uses up to three bytes for each character to move data and performs additional checks at the session level to ensure data integrity.
- PowerCenter repository. The PowerCenter repository can store data in any language. You can use the UTF-8 code page for the PowerCenter repository to store multibyte data in the PowerCenter repository. The code page for the PowerCenter repository is the same as the database code page.
- Metadata Manager repository. The Metadata Manager repository can store data in any language. You can use the UTF-8 code page for the Metadata Manager repository to store multibyte data in the repository. The code page for the repository is the same as the database code page.
- Sources and targets. The sources and targets store data in one or more languages. You use code pages to specify the type of characters in the sources and targets.
- PowerCenter command line programs. You must also ensure that the code page for *pmrep* is a subset of the PowerCenter repository code page and the code page for *pmcmd* is a subset of the PowerCenter Integration Service process code page.

Most database servers use two code pages, a client code page to receive data from client applications and a server code page to store the data. When the database server is running, it converts data between the two code pages if they are different. In this type of database configuration, the PowerCenter Integration Service process interacts with the database client code page. Thus, code pages used by the PowerCenter Integration Service process, such as the PowerCenter repository, source, or target code pages, must be identical to the database client code page. The database client code page is usually identical to the operating system code page on which the PowerCenter Integration Service process runs. The database client code page is a subset of the database server code page.

For more information about specific database client and server code pages, see your database documentation.

Domain Configuration Database Code Page

The domain configuration database must be compatible with the code pages of the PowerCenter repository, Metadata Manager repository, and Model repository.

The Service Manager synchronizes the list of users in the domain with the list of users and groups in each application service. If a user name in the domain has characters that the code page of the application service does not recognize, characters do not convert correctly and inconsistencies occur.

Administrator Tool Code Page

The Administrator tool can run on any node in a Informatica domain. The Administrator tool code page is the code page of the operating system of the node. Each node in the domain must use the same code page.

The Administrator tool code page must be:

- A subset of the PowerCenter repository code page
- A subset of the Metadata Manager repository code page
- A subset of the Model Repository code page

PowerCenter Client Code Page

The PowerCenter Client code page is the code page of the operating system of the PowerCenter Client. To communicate with the PowerCenter repository, the PowerCenter Client code page must be a subset of the PowerCenter repository code page.

PowerCenter Integration Service Process Code Page

The code page of a PowerCenter Integration Service process is the code page of the node that runs the PowerCenter Integration Service process. Define the code page for each PowerCenter Integration Service process in the Administrator tool on the Processes tab.

However, on UNIX, you can change the code page of the PowerCenter Integration Service process by changing the LANG, LC_CTYPE or LC_ALL environment variable for the user that starts the process.

The code page of the PowerCenter Integration Service process must be:

- A subset of the PowerCenter repository code page
- A superset of the machine hosting *pmcmd* or a superset of the code page specified in the INFA_CODEPAGENAME environment variable

The code pages of all PowerCenter Integration Service processes must be compatible with each other. For example, you can use MS Windows Latin1 for a node on Windows and ISO-8859-1 for a node on UNIX.

PowerCenter Integration Services configured for Unicode mode validate code pages when you start a session to ensure accurate data movement. It uses session code pages to convert character data. When the PowerCenter Integration Service runs in ASCII mode, it does not validate session code pages. It reads all character data as ASCII characters and does not perform code page conversions.

Each code page has associated sort orders. When you configure a session, you can select one of the sort orders associated with the code page of the PowerCenter Integration Service process. When you run the PowerCenter Integration Service in Unicode mode, it uses the selected session sort order to sort character data. When you run the PowerCenter Integration Service in ASCII mode, it sorts all character data using a binary sort order.

If you run the PowerCenter Integration Service in the United States on Windows, consider using MS Windows Latin1 (ANSI) as the code page of the PowerCenter Integration Service process.

If you run the PowerCenter Integration Service in the United States on UNIX, consider using ISO 8859-1 as the code page for the PowerCenter Integration Service process.

If you use *pmcmd* to communicate with the PowerCenter Integration Service, the code page of the operating system hosting *pmcmd* must be identical to the code page of the PowerCenter Integration Service process.

The PowerCenter Integration Service generates the names of session log files, reject files, caches and cache files, and performance detail files based on the code page of the PowerCenter Integration Service process.

PowerCenter Repository Code Page

The PowerCenter repository code page is the code page of the data in the repository. The PowerCenter Repository Service uses the PowerCenter repository code page to save metadata in and retrieve metadata from the PowerCenter repository database. Choose the PowerCenter repository code page when you create or upgrade a PowerCenter repository. When the PowerCenter repository database code page is UTF-8, you can create a PowerCenter repository using UTF-8 as its code page.

The PowerCenter repository code page must be:

- Compatible with the domain configuration database code page
- A superset of the the Administrator tool code page
- A superset of the PowerCenter Client code page
- A superset of the code page for the PowerCenter Integration Service process
- A superset of the machine hosting *pmrep* or a superset of the code page specified in the `INFA_CODEPAGENAME` environment variable

A global PowerCenter repository code page must be a subset of the local PowerCenter repository code page if you want to create shortcuts in the local PowerCenter repository that reference an object in a global PowerCenter repository.

If you copy objects from one PowerCenter repository to another PowerCenter repository, the code page for the target PowerCenter repository must be a superset of the code page for the source PowerCenter repository.

Metadata Manager Repository Code Page

The Metadata Manager repository code page is the code page of the data in the repository. The Metadata Manager Service uses the Metadata Manager repository code page to save metadata to and retrieve metadata from the repository database. The Administrator tool writes user and group information to the Metadata Manager Service. The Administrator tool also writes domain information in the repository database. The PowerCenter Integration Service process writes metadata to the repository database. Choose the repository code page when you create or upgrade a Metadata Manager repository. When the repository database code page is UTF-8, you can create a repository using UTF-8 as its code page.

The Metadata Manager repository code page must be:

- Compatible with the domain configuration database code page
- A superset of the Administrator tool code page
- A subset of the PowerCenter repository code page
- A superset of the code page for the PowerCenter Integration Service process

PowerCenter Source Code Page

The source code page depends on the type of source:

- Flat files and VSAM files. The code page of the data in the file. When you configure the flat file or COBOL source definition, choose a code page that matches the code page of the data in the file.
- XML files. The PowerCenter Integration Service converts XML to Unicode when it parses an XML source. When you create an XML source definition, the PowerCenter Designer assigns a default code page. You cannot change the code page.
- Relational databases. The code page of the database client. When you configure the relational connection in the PowerCenter Workflow Manager, choose a code page that is compatible with the code page of the

database client. If you set a database environment variable to specify the language for the database, ensure the code page for the connection is compatible with the language set for the variable. For example, if you set the NLS_LANG environment variable for an Oracle database, ensure that the code page of the Oracle connection is identical to the value set in the NLS_LANG variable. If you do not use compatible code pages, sessions may hang, data may become inconsistent, or you might receive a database error, such as:

```
ORA-00911: Invalid character specified.
```

Regardless of the type of source, the source code page must be a subset of the code page of transformations and targets that receive data from the source. The source code page does not need to be a subset of transformations or targets that do not receive data from the source.

Note: Select IBM EBCDIC as the source database connection code page only if you access EBCDIC data, such as data from a mainframe extract file.

PowerCenter Target Code Page

The target code page depends on the type of target:

- Flat files. When you configure the flat file target definition, choose a code page that matches the code page of the data in the flat file.
- XML files. Configure the XML target code page after you create the XML target definition. The XML Wizard assigns a default code page to the XML target. The PowerCenter Designer does not apply the code page that appears in the XML schema.
- Relational databases. When you configure the relational connection in the PowerCenter Workflow Manager, choose a code page that is compatible with the code page of the database client. If you set a database environment variable to specify the language for the database, ensure the code page for the connection is compatible with the language set for the variable. For example, if you set the NLS_LANG environment variable for an Oracle database, ensure that the code page of the Oracle connection is compatible with the value set in the NLS_LANG variable. If you do not use compatible code pages, sessions may hang or you might receive a database error, such as:

```
ORA-00911: Invalid character specified.
```

The target code page must be a superset of the code page of transformations and sources that provide data to the target. The target code page does not need to be a superset of transformations or sources that do not provide data to the target.

The PowerCenter Integration Service creates session indicator files, session output files, and external loader control and data files using the target flat file code page.

Note: Select IBM EBCDIC as the target database connection code page only if you access EBCDIC data, such as data from a mainframe extract file.

Command Line Program Code Pages

The *pmcmd* and *pmrep* command line programs require code page compatibility. *pmcmd* and *pmrep* use code pages when sending commands in Unicode. Other command line programs do not require code pages.

The code page compatibility for *pmcmd* and *pmrep* depends on whether you configured the code page environment variable INFA_CODEPAGENAME for *pmcmd* or *pmrep*. You can set this variable for either command line program or for both.

If you did not set this variable for a command line program, ensure the following requirements are met:

- If you did not set the variable for *pmcmd*, then the code page of the machine hosting *pmcmd* must be a subset of the code page for the PowerCenter Integration Service process.

- If you did not set the variable for *pmrep*, then the code page of the machine hosting *pmrep* must be a subset of the PowerCenter repository code page.

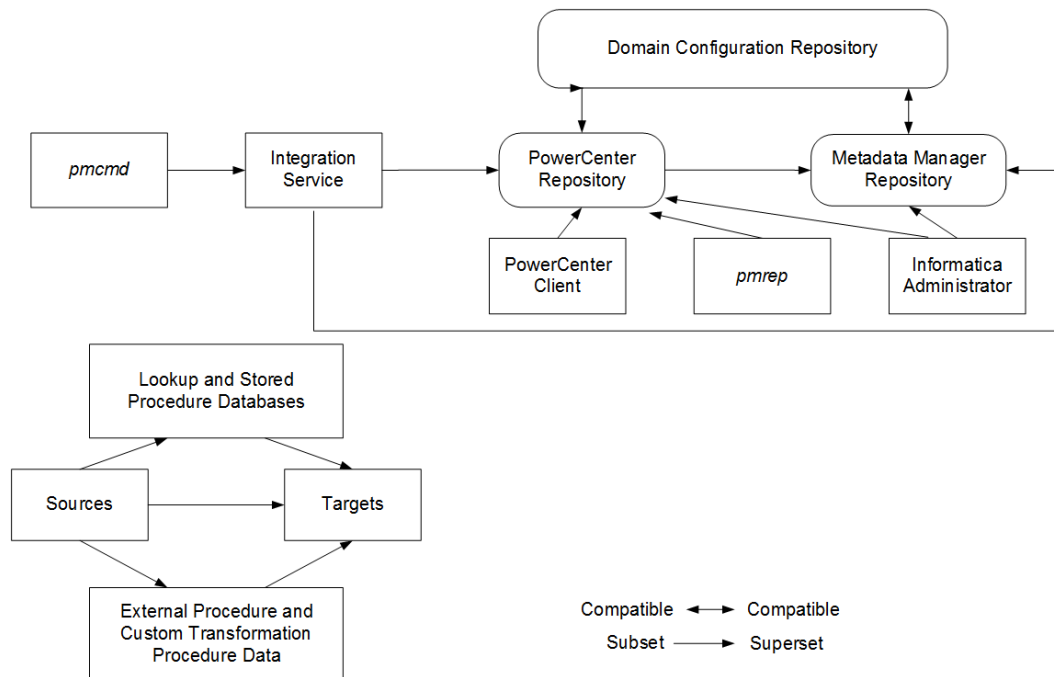
If you set the code page environment variable `INFA_CODEPAGE` for *pmcmd* or *pmrep*, ensure the following requirements are met:

- If you set `INFA_CODEPAGE` for *pmcmd*, the code page defined for the variable must be a subset of the code page for the PowerCenter Integration Service process.
- If you set `INFA_CODEPAGE` for *pmrep*, the code page defined for the variable must be a subset of the PowerCenter repository code page.
- If you run *pmcmd* and *pmrep* from the same machine and you set the `INFA_CODEPAGE` variable, the code page defined for the variable must be subsets of the code pages for the PowerCenter Integration Service process and the PowerCenter repository.

If the code pages are not compatible, the PowerCenter Integration Service process may not fetch the workflow, session, or task from the PowerCenter repository.

Code Page Compatibility Summary

The following image shows code page compatibility in the Informatica environment:



The following table summarizes code page compatibility between sources, targets, repositories, the Informatica Administrator, PowerCenter Client, and Integration Service process:

Component Code Page	Code Page Compatibility
Source (including relational, flat file, and XML file)	Subset of target. Subset of lookup data. Subset of stored procedures. Subset of External Procedure or Custom transformation procedure code page.
Target (including relational, XML files, and flat files)	Superset of source. Superset of lookup data. Superset of stored procedures. Superset of External Procedure or Custom transformation procedure code page. Integration Service process creates external loader data and control files using the target flat file code page.
Lookup and stored procedure database	Subset of target. Superset of source.
External Procedure and Custom transformation procedures	Subset of target. Superset of source.
Domain Configuration Database	Compatible with the PowerCenter Repository Service. Compatible with the Metadata Manager repository.
PowerCenter Integration Service process	Compatible with its operating system. Subset of the PowerCenter repository. Subset of the Metadata Manager repository. Superset of the machine hosting <i>pmcmd</i> . Identical to other nodes running the PowerCenter Integration Service processes.
PowerCenter repository	Compatible with the domain configuration database. Superset of PowerCenter Client. Superset of the nodes running the PowerCenter Integration Service process. Superset of the Metadata Manager repository. A global PowerCenter repository code page must be a subset of a local PowerCenter repository.
PowerCenter Client	Subset of the PowerCenter repository.
Machine running <i>pmcmd</i>	Subset of the PowerCenter Integration Service process.
Machine running <i>pmrep</i>	Subset of the PowerCenter repository.

Component Code Page	Code Page Compatibility
Administrator Tool	Subset of the PowerCenter repository. Subset of the Metadata Manager repository.
Metadata Manager repository	Compatible with the domain configuration database. Subset of the PowerCenter repository. Superset of the Administrator tool. Superset of the PowerCenter Integration Service process.

Code Page Validation

The machines hosting the PowerCenter Client, PowerCenter Integration Service process, and PowerCenter repository database must use appropriate code pages. This eliminates the risk of data or repository inconsistencies. When the PowerCenter Integration Service runs in Unicode data movement mode, it enforces session code page relationships. When the PowerCenter Integration Service runs in ASCII mode, it does not enforce session code page relationships.

To ensure compatibility, the PowerCenter Client and PowerCenter Integration Service perform the following code page validations:

- PowerCenter restricts the use of EBCDIC-based code pages for repositories. Since you cannot install the PowerCenter Client or PowerCenter repository on mainframe systems, you cannot select EBCDIC-based code pages, like IBM EBCDIC, as the PowerCenter repository code page.
- PowerCenter Client can connect to the PowerCenter repository when its code page is a subset of the PowerCenter repository code page. If the PowerCenter Client code page is not a subset of the PowerCenter repository code page, the PowerCenter Client fails to connect to the PowerCenter repository code page with the following error:

```
REP_61082 <PowerCenter Client>'s code page <PowerCenter Client code page> is not one-way compatible to repository <PowerCenter repository name>'s code page <PowerCenter repository code page>.
```

- After you set the PowerCenter repository code page, you cannot change it. After you create or upgrade a PowerCenter repository, you cannot change the PowerCenter repository code page. This prevents data loss and inconsistencies in the PowerCenter repository.
- The PowerCenter Integration Service process can start if its code page is a subset of the PowerCenter repository code page. The code page of the PowerCenter Integration Service process must be a subset of the PowerCenter repository code page to prevent data loss or inconsistencies. If it is not a subset of the PowerCenter repository code page, the PowerCenter Integration Service writes the following message to the log files:

```
REP_61082 <PowerCenter Integration Service>'s code page <PowerCenter Integration Service code page> is not one-way compatible to repository <PowerCenter repository name>'s code page <PowerCenter repository code page>.
```

- When in Unicode data movement mode, the PowerCenter Integration Service starts workflows with the appropriate source and target code page relationships for each session. When the PowerCenter Integration Service runs in Unicode mode, the code page for every source in a session must be a subset of the target code page. This prevents data loss during a session.

If the source and target code pages do not have the appropriate relationships with each other, the PowerCenter Integration Service fails the session and writes the following message to the session log:

```
TM_6227 Error: Code page incompatible in session <session name>. <Additional details>.
```

- The PowerCenter Workflow Manager validates source, target, lookup, and stored procedure code page relationships for each session. The PowerCenter Workflow Manager checks code page relationships when you save a session, regardless of the PowerCenter Integration Service data movement mode. If you configure a session with invalid source, target, lookup, or stored procedure code page relationships, the PowerCenter Workflow Manager issues a warning similar to the following when you save the session:

```
CMN_1933 Code page <code page name> for data from file or connection associated with transformation <name of source, target, or transformation> needs to be one-way compatible with code page <code page name> for transformation <source or target or transformation name>.
```

If you want to run the session in ASCII mode, you can save the session as configured. If you want to run the session in Unicode mode, edit the session to use appropriate code pages.

Relaxed Code Page Validation

Your environment may require you to process data from different sources using character sets from different languages. For example, you may need to process data from English and Japanese sources using the same PowerCenter repository, or you may want to extract source data encoded in a Unicode encoding such as UTF-8. You can configure the PowerCenter Integration Service for relaxed code page validation. Relaxed code page validation enables you to process data using sources and targets with incompatible code pages.

Although relaxed code page validation removes source and target code page restrictions, it still enforces code page compatibility between the PowerCenter Integration Service and PowerCenter repository.

Note: Relaxed code page validation does not safeguard against possible data inconsistencies when you move data between incompatible code pages. You must verify that the characters the PowerCenter Integration Service reads from the source are included in the target code page.

Informatica removes the following restrictions when you relax code page validation:

- Source and target code pages. You can use any code page supported by Informatica for your source and target data.
- Session sort order. You can use any sort order supported by Informatica when you configure a session.

When you run a session with relaxed code page validation, the PowerCenter Integration Service writes the following message to the session log:

```
TM_6185 WARNING! Data code page validation is disabled in this session.
```

When you relax code page validation, the PowerCenter Integration Service writes descriptions of the database connection code pages to the session log.

The following text shows sample code page messages in the session log:

```
TM_6187 Repository code page: [MS Windows Latin 1 (ANSI), superset of Latin 1]
WRT_8222 Target file [$PMTargetFileDir\passthru.out] code page: [MS Windows Traditional Chinese, superset of Big 5]
WRT_8221 Target database connection [Japanese Oracle] code page: [MS Windows Japanese, superset of Shift-JIS]
TM_6189 Source database connection [Japanese Oracle] code page: [MS Windows Japanese, superset of Shift-JIS]
CMN_1716 Lookup [LKP_sjis_lookup] uses database connection [Japanese Oracle] in code page [MS Windows Japanese, superset of Shift-JIS]
CMN_1717 Stored procedure [J_SP_INCREMENT] uses database connection [Japanese Oracle] in code page [MS Windows Japanese, superset of Shift-JIS]
```

If the PowerCenter Integration Service cannot correctly convert data, it writes an error message to the session log.

Configuring the PowerCenter Integration Service

To configure the PowerCenter Integration Service for code page relaxation, complete the following tasks in the Administrator tool:

- Disable code page validation. Disable the `ValidateDataCodePages` option in the PowerCenter Integration Service properties.
- Configure the PowerCenter Integration Service for Unicode data movement mode. Select Unicode for the Data Movement Mode option in the PowerCenter Integration Service properties.
- Configure the PowerCenter Integration Service to write to the logs using the UTF-8 character set. If you configure sessions or workflows to write to log files, enable the `LogInUTF8` option in the PowerCenter Integration Service properties. The PowerCenter Integration Service writes all logs in UTF-8 when you enable the `LogInUTF8` option. The PowerCenter Integration Service writes to the Log Manager in UTF-8 by default.

Selecting Compatible Source and Target Code Pages

Although PowerCenter allows you to use any supported code page, there are risks associated with using incompatible code pages for sources and targets. If your target code page is not a superset of your source code page, you risk inconsistencies in the target data because the source data may contain characters not encoded in the target code page.

When the PowerCenter Integration Service reads characters that are not included in the target code page, you risk transformation errors, inconsistent data, or failed sessions.

Note: If you relax code page validation, it is your responsibility to ensure that data converts from the source to target properly.

Troubleshooting for Code Page Relaxation

The PowerCenter Integration Service failed a session and wrote the following message to the session log:

```
TM_6188 The specified sort order is incompatible with the PowerCenter Integration
Service code page.
```

If you want to validate code pages, select a sort order compatible with the PowerCenter Integration Service code page. If you want to relax code page validation, configure the PowerCenter Integration Service to relax code page validation in Unicode data movement mode.

I tried to view the session or workflow log, but it contains garbage characters.

The PowerCenter Integration Service is not configured to write session or workflow logs using the UTF-8 character set.

Enable the `LogInUTF8` option in the PowerCenter Integration Service properties.

PowerCenter Code Page Conversion

When in data movement mode is set to Unicode, the PowerCenter Client accepts input in any language and converts it to UCS-2. The PowerCenter Integration Service converts source data to UCS-2 before processing and converts the processed data from UCS-2 to the target code page before loading.

When you run a session, the PowerCenter Integration Service converts source, target, and lookup queries from the PowerCenter repository code page to the source, target, or lookup code page. The PowerCenter Integration Service also converts the name and call text of stored procedures from the PowerCenter repository code page to the stored procedure database code page.

At run time, the PowerCenter Integration Service verifies that it can convert the following queries and procedure text from the PowerCenter repository code page without data loss:

- Source query. Must convert to source database code page.
- Lookup query. Must convert to lookup database code page.
- Target SQL query. Must convert to target database code page.
- Name and call text of stored procedures. Must convert to stored procedure database code page.

Choosing Characters for PowerCenter Repository Metadata

You can use any character in the PowerCenter repository code page when inputting PowerCenter repository metadata. If the PowerCenter repository uses UTF-8, you can input any Unicode character. For example, you can store German, Japanese, and English metadata in a UTF-8 enabled PowerCenter repository. However, you must ensure that the PowerCenter Integration Service can successfully perform SQL transactions with source, target, lookup, and stored procedure databases. You must also ensure that the PowerCenter Integration Service can read from source and lookup files and write to target and lookup files. Therefore, when you run a session, you must ensure that the PowerCenter repository metadata characters are encoded in the source, target, lookup, and stored procedure code pages.

Example

The PowerCenter Integration Service, PowerCenter repository, and PowerCenter Client use the ISO 8859-1 Latin1 code page, and the source database contains Japanese data encoded using the Shift-JIS code page. Each code page contains characters not encoded in the other. Using characters other than 7-bit ASCII for the PowerCenter repository and source database metadata can cause the sessions to fail or load no rows to the target in the following situations:

- You create a mapping that contains a string literal with characters specific to the German language range of ISO 8859-1 in a query. The source database may reject the query or return inconsistent results.
- You use the PowerCenter Client to generate SQL queries containing characters specific to the German language range of ISO 8859-1. The source database cannot convert the German-specific characters from the ISO 8859-1 code page into the Shift-JIS code page.
- The source database has a table name that contains Japanese characters. The PowerCenter Designer cannot convert the Japanese characters from the source database code page to the PowerCenter Client code page. Instead, the PowerCenter Designer imports the Japanese characters as question marks (?), changing the name of the table. The PowerCenter Repository Service saves the source table name in the PowerCenter repository as question marks. If the PowerCenter Integration Service sends a query to the source database using the changed table name, the source database cannot find the correct table, and returns no rows or an error to the PowerCenter Integration Service, causing the session to fail.

Because the US-ASCII code page is a subset of both the ISO 8859-1 and Shift-JIS code pages, you can avoid these data inconsistencies if you use 7-bit ASCII characters for all of your metadata.

Case Study: Processing ISO 8859-1 Data

This case study describes how you might set up an environment to process ISO 8859-1 multibyte data. You might want to configure your environment this way if you need to process data from different Western European languages with character sets contained in the ISO 8859-1 code page. This example describes an environment that processes English and German language data.

For this case study, the ISO 8859-1 environment consists of the following elements:

- The PowerCenter Integration Service on a UNIX system
- PowerCenter Client on a Windows system, purchased in the United States
- The PowerCenter repository stored on an Oracle database on UNIX
- A source database containing English language data
- Another source database containing German and English language data
- A target database containing German and English language data
- A lookup database containing English language data

The data environment must process English and German character data.

Configuring the ISO 8859-1 Environment

Use the following guidelines when you configure an environment similar to this case study for ISO 8859-1 data processing:

1. Verify code page compatibility between the PowerCenter repository database client and the database server.
2. Verify code page compatibility between the PowerCenter Client and the PowerCenter repository, and between the PowerCenter Integration Service process and the PowerCenter repository.
3. Set the PowerCenter Integration Service data movement mode to ASCII.
4. Verify session code page compatibility.
5. Verify lookup and stored procedure database code page compatibility.
6. Verify External Procedure or Custom transformation procedure code page compatibility.
7. Configure session sort order.

Step 1. Verify PowerCenter Repository Database Client and Server Compatibility

The database client and server hosting the PowerCenter repository must be able to communicate without data loss.

The PowerCenter repository resides in an Oracle database. Use `NLS_LANG` to set the locale (language, territory, and character set) you want the database client and server to use with your login:

```
NLS_LANG = LANGUAGE_TERRITORY.CHARACTERSET
```

By default, Oracle configures `NLS_LANG` for the U.S. English language, the U.S. territory, and the 7-bit ASCII character set:

```
NLS_LANG = AMERICAN_AMERICA.US7ASCII
```

Change the default configuration to write ISO 8859-1 data to the PowerCenter repository using the Oracle `WE8ISO8859P1` code page. For example:

```
NLS_LANG = AMERICAN_AMERICA.WE8ISO8859P1
```

For more information about verifying and changing the PowerCenter repository database code page, see your database documentation.

Step 2. Verify PowerCenter Code Page Compatibility

The PowerCenter Integration Service and PowerCenter Client code pages must be subsets of the PowerCenter repository code page. Because the PowerCenter Client and PowerCenter Integration Service each use the system code pages of the machines they are installed on, you must verify that the system code pages are subsets of the PowerCenter repository code page.

In this case, the PowerCenter Client on Windows systems were purchased in the United States. Thus the system code pages for the PowerCenter Client machines are set to MS Windows Latin1 by default. To verify system input and display languages, open the Regional Options dialog box from the Windows Control Panel. For systems purchased in the United States, the Regional Settings and Input Locale must be configured for English (United States).

The PowerCenter Integration Service is installed on a UNIX machine. The default code page for UNIX operating systems is ASCII. In this environment, change the UNIX system code page to ISO 8859-1 Western European so that it is a subset of the PowerCenter repository code page.

Step 3. Configure the PowerCenter Integration Service for ASCII Data Movement Mode

Configure the PowerCenter Integration Service to process ISO 8859-1 data. In the Administrator tool, set the Data Movement Mode to ASCII for the PowerCenter Integration Service.

Step 4. Verify Session Code Page Compatibility

When you run a workflow in ASCII data movement mode, the PowerCenter Integration Service enforces source and target code page relationships. To guarantee accurate data conversion, the source code page must be a subset of the target code page.

In this case, the environment contains source databases containing German and English data. When you configure a source database connection in the PowerCenter Workflow Manager, the code page for the connection must be identical to the source database code page and must be a subset of the target code page. Since both the MS Windows Latin1 and the ISO 8859-1 Western European code pages contain German characters, you would most likely use one of these code pages for source database connections.

Because the target code page must be a superset of the source code page, use either MS Windows Latin1, ISO 8859-1 Western European, or UTF-8 for target database connection or flat file code pages. To ensure data consistency, the configured target code page must match the target database or flat file system code page.

If you configure the PowerCenter Integration Service for relaxed code page validation, the PowerCenter Integration Service removes restrictions on source and target code page compatibility. You can select any supported code page for source and target data. However, you must ensure that the targets only receive character data encoded in the target code page.

Step 5. Verify Lookup and Stored Procedure Database Code Page Compatibility

Lookup and stored procedure database code pages must be supersets of the source code pages and subsets of the target code pages. In this case, all lookup and stored procedure database connections must use a code page compatible with the ISO 8859-1 Western European or MS Windows Latin1 code pages.

Step 6. Verify External Procedure or Custom Transformation Procedure Compatibility

External Procedure and Custom transformation procedures must be able to process character data from the source code pages, and they must pass characters that are compatible in the target code pages. In this case, all data processed by the External Procedure or Custom transformations must be in the ISO 8859-1 Western European or MS Windows Latin1 code pages.

Step 7. Configure Session Sort Order

When you run the PowerCenter Integration Service in ASCII mode, it uses a binary sort order for all sessions. In the session properties, the PowerCenter Workflow Manager lists all sort orders associated with the PowerCenter Integration Service code page. You can select a sort order for the session.

Case Study: Processing Unicode UTF-8 Data

This case study describes how you might set up an environment that processes Unicode UTF-8 multibyte data. You might want to configure your environment this way if you need to process data from Western European, Middle Eastern, Asian, or any other language with characters encoded in the UTF-8 character set. This example describes an environment that processes German and Japanese language data.

For this case study, the UTF-8 environment consists of the following elements:

- The PowerCenter Integration Service on a UNIX machine
- The PowerCenter Clients on Windows systems
- The PowerCenter repository stored on an Oracle database on UNIX
- A source database contains German language data
- A source database contains German and Japanese language data
- A target database contains German and Japanese language data
- A lookup database contains German language data

The data environment must process German and Japanese character data.

Configuring the UTF-8 Environment

Use the following guidelines when you configure an environment similar to this case study for UTF-8 data processing:

1. Verify code page compatibility between the PowerCenter repository database client and the database server.
2. Verify code page compatibility between the PowerCenter Client and the PowerCenter repository, and between the PowerCenter Integration Service and the PowerCenter repository.
3. Configure the PowerCenter Integration Service for Unicode data movement mode.
4. Verify session code page compatibility.
5. Verify lookup and stored procedure database code page compatibility.
6. Verify External Procedure or Custom transformation procedure code page compatibility.
7. Configure session sort order.

Step 1. Verify PowerCenter Repository Database Client and Server Code Page Compatibility

The database client and server hosting the PowerCenter repository must be able to communicate without data loss.

The PowerCenter repository resides in an Oracle database. With Oracle, you can use NLS_LANG to set the locale (language, territory, and character set) you want the database client and server to use with your login:

```
NLS_LANG = LANGUAGE_TERRITORY.CHARACTERSET
```

By default, Oracle configures NLS_LANG for U.S. English language, the U.S. territory, and the 7-bit ASCII character set:

```
NLS_LANG = AMERICAN_AMERICA.US7ASCII
```

Change the default configuration to write UTF-8 data to the PowerCenter repository using the Oracle UTF8 character set. For example:

```
NLS_LANG = AMERICAN_AMERICA.UTF8
```

For more information about verifying and changing the PowerCenter repository database code page, see your database documentation.

Step 2. Verify PowerCenter Code Page Compatibility

The PowerCenter Integration Service and PowerCenter Client code pages must be subsets of the PowerCenter repository code page. Because the PowerCenter Client and PowerCenter Integration Service each use the system code pages of the machines they are installed on, you must verify that the system code pages are subsets of the PowerCenter repository code page.

In this case, the PowerCenter Client on Windows systems were purchased in Switzerland. Thus, the system code pages for the PowerCenter Client machines are set to MS Windows Latin1 by default. To verify system input and display languages, open the Regional Options dialog box from the Windows Control Panel.

The PowerCenter Integration Service is installed on a UNIX machine. The default code page for UNIX operating systems is ASCII. In this environment, the UNIX system character set must be changed to UTF-8.

Step 3. Configure the PowerCenter Integration Service for Unicode Data Movement Mode

You must configure the PowerCenter Integration Service to process UTF-8 data. In the Administrator tool, set the Data Movement Mode to Unicode for the PowerCenter Integration Service. The PowerCenter Integration Service allots an extra byte for each character when processing multibyte data.

Step 4. Verify Session Code Page Compatibility

When you run a PowerCenter workflow in Unicode data movement mode, the PowerCenter Integration Service enforces source and target code page relationships. To guarantee accurate data conversion, the source code page must be a subset of the target code page.

In this case, the environment contains a source database containing German and Japanese data. When you configure a source database connection in the PowerCenter Workflow Manager, the code page for the connection must be identical to the source database code page. You can use any code page for the source database.

Because the target code page must be a superset of the source code pages, you must use UTF-8 for the target database connections or flat files. To ensure data consistency, the configured target code page must match the target database or flat file system code page.

If you configure the PowerCenter Integration Service for relaxed code page validation, the PowerCenter Integration Service removes restrictions on source and target code page compatibility. You can select any supported code page for source and target data. However, you must ensure that the targets only receive character data encoded in the target code page.

Step 5. Verify Lookup and Stored Procedure Database Code Page Compatibility

Lookup and stored procedure database code pages must be supersets of the source code pages and subsets of the target code pages. In this case, all lookup and stored procedure database connections must use a code page compatible with UTF-8.

Step 6. Verify External Procedure or Custom Transformation Procedure Compatibility

External Procedure and Custom transformation procedures must be able to process character data from the source code pages, and they must pass characters that are compatible in the target code pages.

In this case, the External Procedure or Custom transformations must be able to process the German and Japanese data from the sources. However, the PowerCenter Integration Service passes data to procedures in UCS-2. Therefore, all data processed by the External Procedure or Custom transformations must be in the UCS-2 character set.

Step 7. Configure Session Sort Order

When you run the PowerCenter Integration Service in Unicode mode, it sorts session data using the sort order configured for the session. By default, sessions are configured for a binary sort order.

To sort German and Japanese data when the PowerCenter Integration Service uses UTF-8, you most likely want to use the default binary sort order.

APPENDIX A

Code Pages

This appendix includes the following topics:

- [Supported Code Pages for Application Services, 323](#)
- [Supported Code Pages for Sources and Targets, 325](#)

Supported Code Pages for Application Services

Informatica supports code pages for internationalization. Informatica uses International Components for Unicode (ICU) for its globalization support. For a list of code page aliases in ICU, see <http://demo.icu-project.org/icu-bin/convexp>.

When you assign an application service code page in the Administrator tool, you select the code page description.

You must use UTF-8 compatible code pages for the domain, Model Repository Service, and for each Data Integration Service process.

The following table lists the name, description, and ID for supported code pages for the PowerCenter Repository Service, the Metadata Manager Service, and for each PowerCenter Integration Service process:

Name	Description	ID
IBM037	IBM EBCDIC US English	2028
IBM1047	IBM EBCDIC US English IBM1047	1047
IBM273	IBM EBCDIC German	2030
IBM280	IBM EBCDIC Italian	2035
IBM285	IBM EBCDIC UK English	2038
IBM297	IBM EBCDIC French	2040
IBM500	IBM EBCDIC International Latin-1	2044
IBM930	IBM EBCDIC Japanese	930
IBM935	IBM EBCDIC Simplified Chinese	935

Name	Description	ID
IBM937	IBM EBCDIC Traditional Chinese	937
IBM939	IBM EBCDIC Japanese CP939	939
ISO-8859-10	ISO 8859-10 Latin 6 (Nordic)	13
ISO-8859-15	ISO 8859-15 Latin 9 (Western European)	201
ISO-8859-2	ISO 8859-2 Eastern European	5
ISO-8859-3	ISO 8859-3 Southeast European	6
ISO-8859-4	ISO 8859-4 Baltic	7
ISO-8859-5	ISO 8859-5 Cyrillic	8
ISO-8859-6	ISO 8859-6 Arabic	9
ISO-8859-7	ISO 8859-7 Greek	10
ISO-8859-8	ISO 8859-8 Hebrew	11
ISO-8859-9	ISO 8859-9 Latin 5 (Turkish)	12
JapanEUC	Japanese Extended UNIX Code (including JIS X 0212)	18
Latin1	ISO 8859-1 Western European	4
MS1250	MS Windows Latin 2 (Central Europe)	2250
MS1251	MS Windows Cyrillic (Slavic)	2251
MS1252	MS Windows Latin 1 (ANSI), superset of Latin1	2252
MS1253	MS Windows Greek	2253
MS1254	MS Windows Latin 5 (Turkish), superset of ISO 8859-9	2254
MS1255	MS Windows Hebrew	2255
MS1256	MS Windows Arabic	2256
MS1257	MS Windows Baltic Rim	2257
MS1258	MS Windows Vietnamese	2258
MS1361	MS Windows Korean (Johab)	1361
MS874	MS-DOS Thai, superset of TIS 620	874
MS932	MS Windows Japanese, Shift-JIS	2024
MS936	MS Windows Simplified Chinese, superset of GB 2312-80, EUC encoding	936

Name	Description	ID
MS949	MS Windows Korean, superset of KS C 5601-1992	949
MS950	MS Windows Traditional Chinese, superset of Big 5	950
US-ASCII	7-bit ASCII	1
UTF-8	UTF-8 encoding of Unicode	106

Supported Code Pages for Sources and Targets

Informatica supports code pages for internationalization. Informatica uses International Components for Unicode (ICU) for its globalization support. For a list of code page aliases in ICU, see <http://demo.icu-project.org/icu-bin/convexp>.

When you assign a source or target code page in the PowerCenter Client, you select the code page description. When you assign a code page using the *pmrep* CreateConnection command or define a code page in a parameter file, you enter the code page name. The following table lists the name, description, and ID for supported code pages for sources and targets:

Name	Description	ID
Adobe-Standard-Encoding	Adobe Standard Encoding	10073
BOCU-1	Binary Ordered Compression for Unicode (BOCU-1)	10010
CESU-8	ICompatibility Encoding Scheme for UTF-16 (CESU-8)	10011
cp1006	ISO Urdu	10075
cp1098	PC Farsi	10076
cp1124	ISO Cyrillic Ukraine	10077
cp1125	PC Cyrillic Ukraine	10078
cp1131	PC Cyrillic Belarus	10080
cp1381	PC Chinese GB (S-Ch Data mixed)	10082
cp850	PC Latin1	10036
cp851	PC DOS Greek (without euro)	10037
cp856	PC Hebrew (old)	10040
cp857	PC Latin5 (without euro update)	10041
cp858	PC Latin1 (with euro update)	10042

Name	Description	ID
cp860	PC Portugal	10043
cp861	PC Iceland	10044
cp862	PC Hebrew (without euro update)	10045
cp863	PC Canadian French	10046
cp864	PC Arabic (without euro update)	10047
cp865	PC Nordic	10048
cp866	PC Russian (without euro update)	10049
cp868	PC Urdu	10051
cp869	PC Greek (without euro update)	10052
cp922	IPC Estonian (without euro update)	10056
cp949c	PC Korea - KS	10028
ebcdic-xml-us	EBCDIC US (with euro) - Extension for XML4C(Xerces)	10180
EUC-KR	EUC Korean	10029
GB_2312-80	Simplified Chinese (GB2312-80)	10025
gb18030	GB 18030 MBCS codepage	1392
GB2312	Chinese EUC	10024
HKSCS	Hong Kong Supplementary Character Set	9200
hp-roman8	HP Latin1	10072
HZ-GB-2312	Simplified Chinese (HZ GB2312)	10092
IBM037	IBM EBCDIC US English	2028
IBM-1025	EBCDIC Cyrillic	10127
IBM1026	EBCDIC Turkey	10128
IBM1047	IBM EBCDIC US English IBM1047	1047
IBM-1047-s390	EBCDIC IBM-1047 for S/390 (If and nl swapped)	10167
IBM-1097	EBCDIC Farsi	10129
IBM-1112	EBCDIC Baltic	10130
IBM-1122	EBCDIC Estonia	10131

Name	Description	ID
IBM-1123	EBCDIC Cyrillic Ukraine	10132
IBM-1129	ISO Vietnamese	10079
IBM-1130	EBCDIC Vietnamese	10133
IBM-1132	EBCDIC Lao	10134
IBM-1133	ISO Lao	10081
IBM-1137	EBCDIC Devanagari	10163
IBM-1140	EBCDIC US (with euro update)	10135
IBM-1140-s390	EBCDIC IBM-1140 for S/390 (If and nl swapped)	10168
IBM-1141	EBCDIC Germany, Austria (with euro update)	10136
IBM-1142	EBCDIC Denmark, Norway (with euro update)	10137
IBM-1142-s390	EBCDIC IBM-1142 for S/390 (If and nl swapped)	10169
IBM-1143	EBCDIC Finland, Sweden (with euro update)	10138
IBM-1143-s390	EBCDIC IBM-1143 for S/390 (If and nl swapped)	10170
IBM-1144	EBCDIC Italy (with euro update)	10139
IBM-1144-s390	EBCDIC IBM-1144 for S/390 (If and nl swapped)	10171
IBM-1145	EBCDIC Spain, Latin America (with euro update)	10140
IBM-1145-s390	EBCDIC IBM-1145 for S/390 (If and nl swapped)	10172
IBM-1146	EBCDIC UK, Ireland (with euro update)	10141
IBM-1146-s390	EBCDIC IBM-1146 for S/390 (If and nl swapped)	10173
IBM-1147	EBCDIC French (with euro update)	10142
IBM-1147-s390	EBCDIC IBM-1147 for S/390 (If and nl swapped)	10174
IBM-1147-s390	EBCDIC IBM-1147 for S/390 (If and nl swapped)	10174
IBM-1148	EBCDIC International Latin1 (with euro update)	10143
IBM-1148-s390	EBCDIC IBM-1148 for S/390 (If and nl swapped)	10175
IBM-1149	EBCDIC Iceland (with euro update)	10144
IBM-1149-s390	IEBCDIC IBM-1149 for S/390 (If and nl swapped)	10176
IBM-1153	EBCDIC Latin2 (with euro update)	10145

Name	Description	ID
IBM-1153-s390	EBCDIC IBM-1153 for S/390 (If and nl swapped)	10177
IBM-1154	EBCDIC Cyrillic Multilingual (with euro update)	10146
IBM-1155	EBCDIC Turkey (with euro update)	10147
IBM-1156	EBCDIC Baltic Multilingual (with euro update)	10148
IBM-1157	EBCDIC Estonia (with euro update)	10149
IBM-1158	EBCDIC Cyrillic Ukraine (with euro update)	10150
IBM1159	IBM EBCDIC Taiwan, Traditional Chinese	11001
IBM-1160	EBCDIC Thai (with euro update)	10151
IBM-1162	Thai (with euro update)	10033
IBM-1164	EBCDIC Vietnamese (with euro update)	10152
IBM-1250	MS Windows Latin2 (without euro update)	10058
IBM-1251	MS Windows Cyrillic (without euro update)	10059
IBM-1255	MS Windows Hebrew (without euro update)	10060
IBM-1256	MS Windows Arabic (without euro update)	10062
IBM-1257	MS Windows Baltic (without euro update)	10064
IBM-1258	MS Windows Vietnamese (without euro update)	10066
IBM-12712	EBCDIC Hebrew (updated with euro and new sheqel, control characters)	10161
IBM-12712-s390	EBCDIC IBM-12712 for S/390 (If and nl swapped)	10178
IBM-1277	Adobe Latin1 Encoding	10074
IBM13121	IBM EBCDIC Korean Extended CP13121	11002
IBM13124	IBM EBCDIC Simplified Chinese CP13124	11003
IBM-1363	PC Korean KSC MBCS Extended (with \ <-> Won mapping)	10032
IBM-1364	EBCDIC Korean Extended (SBCS IBM-13121 combined with DBCS IBM-4930)	10153
IBM-1371	EBCDIC Taiwan Extended (SBCS IBM-1159 combined with DBCS IBM-9027)	10154
IBM-1373	Taiwan Big-5 (with euro update)	10019

Name	Description	ID
IBM-1375	MS Taiwan Big-5 with HKSCS extensions	10022
IBM-1386	PC Chinese GBK (IBM-1386)	10023
IBM-1388	EBCDIC Chinese GB (S-Ch DBCS-Host Data)	10155
IBM-1390	EBCDIC Japanese Katakana (with euro)	10156
IBM-1399	EBCDIC Japanese Latin-Kanji (with euro)	10157
IBM-16684	EBCDIC Japanese Extended (DBCS IBM-1390 combined with DBCS IBM-1399)	10158
IBM-16804	EBCDIC Arabic (with euro update)	10162
IBM-16804-s390	EBCDIC IBM-16804 for S/390 (lf and nl swapped)	10179
IBM-25546	ISO-2022 encoding for Korean (extension 1)	10089
IBM273	IBM EBCDIC German	2030
IBM277	EBCDIC Denmark, Norway	10115
IBM278	EBCDIC Finland, Sweden	10116
IBM280	IBM EBCDIC Italian	2035
IBM284	EBCDIC Spain, Latin America	10117
IBM285	IBM EBCDIC UK English	2038
IBM290	EBCDIC Japanese Katakana SBCS	10118
IBM297	IBM EBCDIC French	2040
IBM-33722	Japanese EUC (with \ <-> Yen mapping)	10017
IBM367	IBM367	10012
IBM-37-s390	EBCDIC IBM-37 for S/390 (lf and nl swapped)	10166
IBM420	EBCDIC Arabic	10119
IBM424	EBCDIC Hebrew (updated with new sheqel, control characters)	10120
IBM437	PC United States	10035
IBM-4899	EBCDIC Hebrew (with euro)	10159
IBM-4909	ISO Greek (with euro update)	10057
IBM4933	IBM Simplified Chinese CP4933	11004

Name	Description	ID
IBM-4971	EBCDIC Greek (with euro update)	10160
IBM500	IBM EBCDIC International Latin-1	2044
IBM-5050	Japanese EUC (Packed Format)	10018
IBM-5123	EBCDIC Japanese Latin (with euro update)	10164
IBM-5351	MS Windows Hebrew (older version)	10061
IBM-5352	MS Windows Arabic (older version)	10063
IBM-5353	MS Windows Baltic (older version)	10065
IBM-803	EBCDIC Hebrew	10121
IBM833	IBM EBCDIC Korean CP833	833
IBM834	IBM EBCDIC Korean CP834	834
IBM835	IBM Taiwan, Traditional Chinese CP835	11005
IBM836	IBM EBCDIC Simplified Chinese Extended	11006
IBM837	IBM Simplified Chinese CP837	11007
IBM-838	EBCDIC Thai	10122
IBM-8482	EBCDIC Japanese Katakana SBCS (with euro update)	10165
IBM852	PC Latin2 (without euro update)	10038
IBM855	PC Cyrillic (without euro update)	10039
IBM-867	PC Hebrew (with euro update)	10050
IBM870	EBCDIC Latin2	10123
IBM871	EBCDIC Iceland	10124
IBM-874	PC Thai (without euro update)	10034
IBM-875	EBCDIC Greek	10125
IBM-901	PC Baltic (with euro update)	10054
IBM-902	PC Estonian (with euro update)	10055
IBM918	EBCDIC Urdu	10126
IBM930	IBM EBCDIC Japanese	930
IBM933	IBM EBCDIC Korean CP933	933

Name	Description	ID
IBM935	IBM EBCDIC Simplified Chinese	935
IBM937	IBM EBCDIC Traditional Chinese	937
IBM939	IBM EBCDIC Japanese CP939	939
IBM-942	PC Japanese SJIS-78 syntax (IBM-942)	10015
IBM-943	PC Japanese SJIS-90 (IBM-943)	10016
IBM-949	PC Korea - KS (default)	10027
IBM-950	Taiwan Big-5 (without euro update)	10020
IBM-964	EUC Taiwan	10026
IBM-971	EUC Korean (DBCS-only)	10030
IMAP-mailbox-name	IMAP Mailbox Name	10008
is-960	Israeli Standard 960 (7-bit Hebrew encoding)	11000
ISO-2022-CN	ISO-2022 encoding for Chinese	10090
ISO-2022-CN-EXT	ISO-2022 encoding for Chinese (extension 1)	10091
ISO-2022-JP	ISO-2022 encoding for Japanese	10083
ISO-2022-JP-2	ISO-2022 encoding for Japanese (extension 2)	10085
ISO-2022-KR	ISO-2022 encoding for Korean	10088
ISO-8859-10	ISO 8859-10 Latin 6 (Nordic)	13
ISO-8859-13	ISO 8859-13 PC Baltic (without euro update)	10014
ISO-8859-15	ISO 8859-15 Latin 9 (Western European)	201
ISO-8859-2	ISO 8859-2 Eastern European	5
ISO-8859-3	ISO 8859-3 Southeast European	6
ISO-8859-4	ISO 8859-4 Baltic	7
ISO-8859-5	ISO 8859-5 Cyrillic	8
ISO-8859-6	ISO 8859-6 Arabic	9
ISO-8859-7	ISO 8859-7 Greek	10
ISO-8859-8	ISO 8859-8 Hebrew	11
ISO-8859-9	ISO 8859-9 Latin 5 (Turkish)	12

Name	Description	ID
JapanEUC	Japanese Extended UNIX Code (including JIS X 0212)	18
JEF	Japanese EBCDIC Fujitsu	9000
JEF-K	Japanese EBCDIC-Kana Fujitsu	9005
JIPSE	NEC ACOS JIPSE Japanese	9002
JIPSE-K	NEC ACOS JIPSE-Kana Japanese	9007
JIS_Encoding	ISO-2022 encoding for Japanese (extension 1)	10084
JIS_X0201	ISO-2022 encoding for Japanese (JIS_X0201)	10093
JIS7	ISO-2022 encoding for Japanese (extension 3)	10086
JIS8	ISO-2022 encoding for Japanese (extension 4)	10087
JP-EBCDIC	EBCDIC Japanese	9010
JP-EBCDIK	EBCDIK Japanese	9011
KEIS	HITACHI KEIS Japanese	9001
KEIS-K	HITACHI KEIS-Kana Japanese	9006
KOI8-R	IRussian Internet	10053
KSC_5601	PC Korean KSC MBCS Extended (KSC_5601)	10031
Latin1	ISO 8859-1 Western European	4
LMBCS-1	Lotus MBCS encoding for PC Latin1	10103
LMBCS-11	Lotus MBCS encoding for MS-DOS Thai	10110
LMBCS-16	Lotus MBCS encoding for Windows Japanese	10111
LMBCS-17	Lotus MBCS encoding for Windows Korean	10112
LMBCS-18	Lotus MBCS encoding for Windows Chinese (Traditional)	10113
LMBCS-19	Lotus MBCS encoding for Windows Chinese (Simplified)	10114
LMBCS-2	Lotus MBCS encoding for PC DOS Greek	10104
LMBCS-3	Lotus MBCS encoding for Windows Hebrew	10105
LMBCS-4	Lotus MBCS encoding for Windows Arabic	10106
LMBCS-5	Lotus MBCS encoding for Windows Cyrillic	10107
LMBCS-6	Lotus MBCS encoding for PC Latin2	10108

Name	Description	ID
LMBCS-8	Lotus MBCS encoding for Windows Turkish	10109
macintosh	Apple Latin 1	10067
MELCOM	mitsubishi MELCOM Japanese	9004
MELCOM-K	mitsubishi MELCOM-Kana Japanese	9009
MS1250	MS Windows Latin 2 (Central Europe)	2250
MS1251	MS Windows Cyrillic (Slavic)	2251
MS1252	MS Windows Latin 1 (ANSI), superset of Latin1	2252
MS1253	MS Windows Greek	2253
MS1254	MS Windows Latin 5 (Turkish), superset of ISO 8859-9	2254
MS1255	MS Windows Hebrew	2255
MS1256	MS Windows Arabic	2256
MS1257	MS Windows Baltic Rim	2257
MS1258	MS Windows Vietnamese	2258
MS1361	MS Windows Korean (Johab)	1361
MS874	MS-DOS Thai, superset of TIS 620	874
MS932	MS Windows Japanese, Shift-JIS	2024
MS936	MS Windows Simplified Chinese, superset of GB 2312-80, EUC encoding	936
MS949	MS Windows Korean, superset of KS C 5601-1992	949
MS950	MS Windows Traditional Chinese, superset of Big 5	950
SCSU	Standard Compression Scheme for Unicode (SCSU)	10009
UNISYS	UNISYS Japanese	9003
UNISYS-K	UNISYS-Kana Japanese	9008
US-ASCII	7-bit ASCII	1
UTF-16_OppositeEndian	UTF-16 encoding of Unicode (Opposite Platform Endian)	10004
UTF-16_PlatformEndian	UTF-16 encoding of Unicode (Platform Endian)	10003
UTF-16BE	UTF-16 encoding of Unicode (Big Endian)	1200

Name	Description	ID
UTF-16LE	UTF-16 encoding of Unicode (Lower Endian)	1201
UTF-32_OppositeEndian	UTF-32 encoding of Unicode (Opposite Platform Endian)	10006
UTF-32_PlatformEndian	UTF-32 encoding of Unicode (Platform Endian)	10005
UTF-32BE	UTF-32 encoding of Unicode (Big Endian)	10001
UTF-32LE	UTF-32 encoding of Unicode (Lower Endian)	10002
UTF-7	UTF-7 encoding of Unicode	10007
UTF-8	UTF-8 encoding of Unicode	106
windows-57002	Indian Script Code for Information Interchange - Devanagari	10094
windows-57003	Indian Script Code for Information Interchange - Bengali	10095
windows-57004	Indian Script Code for Information Interchange - Tamil	10099
windows-57005	Indian Script Code for Information Interchange - Telugu	10100
windows-57007	Indian Script Code for Information Interchange - Oriya	10098
windows-57008	Indian Script Code for Information Interchange - Kannada	10101
windows-57009	Indian Script Code for Information Interchange - Malayalam	10102
windows-57010	Indian Script Code for Information Interchange - Gujarati	10097
windows-57011	Indian Script Code for Information Interchange - Gurmukhi	10096
x-mac-centraleurroman	Apple Central Europe	10070
x-mac-cyrillic	Apple Cyrillic	10069
x-mac-greek	Apple Greek	10068
x-mac-turkish	Apple Turkish	10071

Restrictions for Code Pages for Sources and Targets

Consider the following restrictions when you assign a source or target code page:

- Select IBM EBCDIC as your source database connection code page only if you access EBCDIC data, such as data from a mainframe extract file.
- The following code pages are not supported for database or relational connections:
 - UTF-16 encoding of Unicode (Opposite Platform Endian)
 - UTF-16 encoding of Unicode (Platform Endian)
 - UTF-16 encoding of Unicode (Big Endian)
 - UTF-16 encoding of Unicode (Lower Endian)

APPENDIX B

Custom Roles

This appendix includes the following topics:

- [Analyst Service Custom Role, 335](#)
- [Metadata Manager Service Custom Roles, 336](#)
- [Operator Custom Role, 337](#)
- [PowerCenter Repository Service Custom Roles, 338](#)
- [Test Data Manager Custom Roles, 339](#)

Analyst Service Custom Role

The Analyst Service Business Glossary Consumer is a custom Analyst Service role.

The following table lists the default privilege assigned to the Analyst Service Business Glossary Consumer custom role:

Privilege Group	Privilege Name
Workspace Access	Glossary Workspace

Metadata Manager Service Custom Roles

Metadata Manager Service custom roles include the Metadata Manager Advanced User, Metadata Manager Basic User, and Metadata Manager Intermediate User roles.

Metadata Manager Advanced User

The following table lists the default privileges assigned to the Metadata Manager Advanced User custom role:

Privilege Group	Privilege Name
Catalog	<ul style="list-style-type: none"> - Share Shortcuts - View Lineage - View Related Catalogs - View Reports - View Profile Results - View Catalog - View Relationships - Manage Relationships - View Comments - Post Comments - Delete Comments - View Links - Manage Links - View Glossary - Manage Objects
Load	<ul style="list-style-type: none"> - View Resource - Load Resource - Manage Schedules - Purge Metadata - Manage Resource
Model	<ul style="list-style-type: none"> - View Model - Manage Model - Export/Import Models
Security	Manage Catalog Permissions

Metadata Manager Basic User

The following table lists the default privileges assigned to the Metadata Manager Basic User custom role:

Privilege Group	Privilege Name
Catalog	<ul style="list-style-type: none"> - View Lineage - View Related Catalogs - View Catalog - View Relationships - View Comments - View Links
Model	View Model

Metadata Manager Intermediate User

The following table lists the default privileges assigned to the Metadata Manager Intermediate User custom role:

Privilege Group	Privilege Name
Catalog	<ul style="list-style-type: none">- View Lineage- View Related Catalogs- View Reports- View Profile Results- View Catalog- View Relationships- View Comments- Post Comments- Delete Comments- View Links- Manage Links- View Glossary
Load	<ul style="list-style-type: none">- View Resource- Load Resource
Model	View Model

Operator Custom Role

The Operator custom role includes privileges for managing, scheduling, and monitoring application services.

The following table lists the default privileges assigned to the Operator custom role:

Privilege Group	Privilege Name
Application Administration	Manage Applications
Domain Administration	Manage Service Execution
Model Repository Service Administration	Manage Team-based Development
Monitoring	<p>The Monitoring privilege group includes the following privileges:</p> <ul style="list-style-type: none">- View: View Jobs of Other Users- View: View Statistics- View: View Reports- Access Monitoring: Access from Analyst Tool- Access Monitoring: Access from Developer Tool- Access Monitoring: Access from Administrator Tool- Perform Actions on Jobs <p>Note: In a domain that uses Kerberos authentication, users must also have the Administrator role for the Model Repository Service that is configured for monitoring.</p>

Privilege Group	Privilege Name
Scheduler	The Scheduler privilege group includes the following privileges: <ul style="list-style-type: none"> - Manage Scheduled Jobs: Create Schedule - Manage Scheduled Jobs: Delete Schedule - Manage Scheduled Jobs: Edit Schedule - Manage Scheduled Jobs: View Schedules
Tools	Access Informatica Administrator

PowerCenter Repository Service Custom Roles

The PowerCenter Repository Service custom roles include the PowerCenter Connection Administrator, PowerCenter Developer, PowerCenter Operator, and PowerCenter Repository Folder Administrator.

PowerCenter Connection Administrator

The following table lists the default privileges assigned to the PowerCenter Connection Administrator custom role:

Privilege Group	Privilege Name
Tools	Access Workflow Manager
Global Objects	Create Connections

PowerCenter Developer

The following table lists the default privileges assigned to the PowerCenter Developer custom role:

Privilege Group	Privilege Name
Tools	<ul style="list-style-type: none"> - Access Designer - Access Workflow Manager - Access Workflow Monitor
Design Objects	<ul style="list-style-type: none"> - Create, Edit, and Delete - Manage Versions
Sources and Targets	<ul style="list-style-type: none"> - Create, Edit, and Delete - Manage Versions
Run-time Objects	<ul style="list-style-type: none"> - Create, Edit, and Delete - Execute - Manage Versions - Monitor

PowerCenter Operator

The following table lists the default privileges assigned to the PowerCenter Operator custom role:

Privilege Group	Privilege Name
Tools	Access Workflow Monitor
Run-time Objects	<ul style="list-style-type: none">- Execute- Manage Execution- Monitor

PowerCenter Repository Folder Administrator

The following table lists the default privileges assigned to the PowerCenter Repository Folder Administrator custom role:

Privilege Group	Privilege Name
Tools	Access Repository Manager
Folders	<ul style="list-style-type: none">- Copy- Create- Manage Versions
Global Objects	<ul style="list-style-type: none">- Manage Deployment Groups- Execute Deployment Groups- Create Labels- Create Queries

Test Data Manager Custom Roles

The Test Data Manager custom roles include the Test Data Administrator, Test Data Developer, Test Data Project DBA, Test Data Project Developer, Test Data Project Owner, Test Data Risk Manager, Test Data Specialist, and Test Engineer.

Test Data Administrator

The following table lists the default privileges assigned to the Test Data Administrator custom role:

Privilege Group	Privilege Name
Projects	Audit Project
Administration	<ul style="list-style-type: none">- View Connections- Manage Connections- Manage Preferences

Test Data Developer

The following table lists the default privileges assigned to the Test Data Developer custom role:

Privilege Group	Privilege Name
Policies	<ul style="list-style-type: none">- View Policies- Manage Policies
Data Domains	<ul style="list-style-type: none">- View Data Domains- Manage Data Domains
Projects	Audit Project

Test Data Project DBA

The following table lists the default privileges assigned to the Test Data Project DBA custom role:

Privilege Group	Privilege Name
Projects	<ul style="list-style-type: none">- View Project- Execute Project- Monitor Project- Audit Project
Administration	<ul style="list-style-type: none">- View Connections- Manage Connections

Test Data Project Developer

The following table lists the default privileges assigned to the Test Data Project Developer custom role:

Privilege Group	Privilege Name
Policies	View Policies
Data Domains	View Data Domains
Projects	<ul style="list-style-type: none">- View Project- Discover Project- Execute Project- Monitor Project- Audit Project- Import Metadata
Data Masking	<ul style="list-style-type: none">- View Data Masking- Manage Data Masking
Data Subset	<ul style="list-style-type: none">- View Data Subset- Manage Data Subset
Administration	<ul style="list-style-type: none">- View Connections- Manage Connections

Test Data Project Owner

The following table lists the default privileges assigned to the Test Data Project Owner custom role:

Privilege Group	Privilege Name
Policies	View Policies
Data Domains	View Data Domains
Projects	<ul style="list-style-type: none">- View Project- Manage Project- Discover Project- Execute Project- Monitor Project- Audit Project- Import Metadata
Data Masking	<ul style="list-style-type: none">- View Data Masking- Manage Data Masking
Data Subset	<ul style="list-style-type: none">- View Data Subset- Manage Data Subset
Administration	<ul style="list-style-type: none">- View Connections- Manage Connections

Test Data Risk Manager

The following table lists the default privileges assigned to the Test Data Risk Manager custom role:

Privilege Group	Privilege Name
Policies	View Policies
Data Domains	View Data Domains
Projects	Audit Project

Test Data Specialist

The following table lists the default privileges assigned to the Test Data Specialist custom role:

Privilege Group	Privilege Name
Policies	View Policies
Data Domains	<ul style="list-style-type: none">- View Data Domains- Manage Data Domains

Privilege Group	Privilege Name
Projects	<ul style="list-style-type: none"> - View Project - Manage Project - Discover Project - Execute Project - Monitor Project - Audit Project - Import Metadata
Data Masking	<ul style="list-style-type: none"> - View Data Masking - Manage Data Masking
Data Subset	<ul style="list-style-type: none"> - View Data Subset - Manage Data Subset
Administration	<ul style="list-style-type: none"> - View Connections - Manage Connections

Test Engineer

The following table lists the default privileges assigned to the Test Engineer custom role:

Privilege Group	Privilege Name
Projects	<ul style="list-style-type: none"> - View Project - Monitor Project

APPENDIX C

Informatica Platform Connectivity

This appendix includes the following topics:

- [Informatica Platform Connectivity Overview, 343](#)
- [Domain Connectivity, 344](#)
- [PowerCenter Connectivity, 346](#)
- [Native Connectivity, 351](#)
- [ODBC Connectivity, 351](#)
- [JDBC Connectivity, 352](#)

Informatica Platform Connectivity Overview

The Informatica platform uses the following types of connectivity to communicate among clients, services, and other components in the domain:

TCP/IP network protocol

Application services and the Service Managers in a domain use TCP/IP network protocol to communicate with other nodes and services. The clients also use TCP/IP to communicate with application services. You can configure the host name and port number for TCP/IP communication on a node when you install the Informatica services. You can configure the port numbers used for services on a node during installation or in Informatica Administrator.

Native drivers

The Data Integration Service uses native drivers to communicate with databases. The PowerCenter Integration Service and the PowerCenter Repository Service use native drivers to communicate with databases. Native drivers are packaged with the database server and client software. Install and configure the native database client software on the machines where the services run.

ODBC

The ODBC drivers are installed with the Informatica services and the Informatica clients. The integration services use ODBC drivers to communicate with databases.

JDBC

The Model Repository Service uses JDBC to connect to the Model repository database. The Metadata Manager Service uses JDBC to connect to the Metadata Manager repository and metadata source repositories.

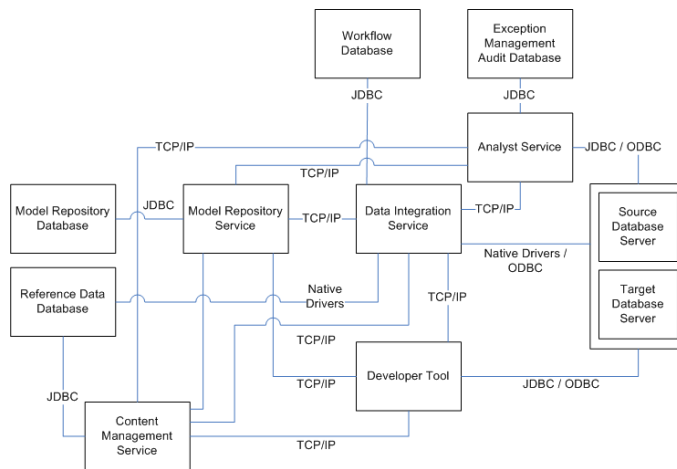
The gateway nodes in the Informatica domain use JDBC to connect to the domain configuration repository.

Domain Connectivity

Services on a node in an Informatica domain use TCP/IP to connect to services on other nodes. Because services can run on multiple nodes in the domain, services rely on the Service Manager to route requests. The Service Manager on the master gateway node handles requests for services and responds with the address of the requested service.

Nodes communicate through TCP/IP on the port you select for a node when you install Informatica Services. When you create a node, you select a port number for the node. The Service Manager listens for incoming TCP/IP connections on that port.

The following figure shows an overview of the connectivity for components in the platform:



The platform uses connection objects to define connectivity information for source and target databases. The connection objects can use native or ODBC connectivity. The Data Integration Service uses connection objects to connect to sources and targets.

The services and clients connect in the following ways:

Model Repository Service

The Model Repository Service uses JDBC to read or write data and metadata in the Model repository. It uses TCP/IP to communicate with the Data Integration Service and the clients.

Data Integration Service

The Data Integration Service uses ODBC or native drivers to connect and read data from a source database and write data to a target database. It uses TCP/IP to communicate with the Model Repository Service, Content Management Service, and client applications.

Informatica Developer

The Developer tool uses TCP/IP to send data transformation requests to the Data Integration Service. It uses TCP/IP to communicate with the Content Management Service to manage reference tables, probabilistic model files, and to retrieve configuration and status information for identity population files and address validation reference data files. When you preview mappings or data objects in the Developer tool, it uses JDBC or ODBC drivers to connect to the source or target database to fetch the metadata required for preview.

Informatica Analyst

The Analyst Service uses TCP/IP to send requests to the Data Integration Service. It uses TCP/IP to communicate with the Content Management Service to manage reference tables. When an Analyst tool user previews profiles or objects, the Analyst Service fetches the metadata that the preview requires

from the source or target database. The Analyst Service uses JDBC or ODBC drivers to connect to the source or target database.

If you use ODBC to connect to the source or target database, install the ODBC driver on the node where the Analyst Service runs.

The Analyst Service can also connect to an exception management audit database. The exception management audit database is a centralized audit trail for the work that Analyst tool users perform on Human task instances. The Analyst Service uses JDBC drivers to connect to the exception management audit database.

Content Management Service


The Content Management Service manages the locations and other properties for reference data. The Content Management Service uses TCP/IP to communicate with the Data Integration Service to read and write data in reference tables. The Content Management Service uses JDBC to communicate directly with the reference data warehouse when it creates reference tables.

If multiple instances of a Content Management Service exist in an Informatica domain, the master Content Management Service updates the Data Integration Service. The master Content Management Service uses TCP/IP to communicate with the Domain Service to identify the Model Repository Service and the Data Integration Service to use.

Model Repository Connectivity

The Model Repository Service connects to the Model repository using JDBC drivers. Informatica Developer, Informatica Analyst, Informatica Administrator, and the Data Integration Service communicate with the Model Repository Service over TCP/IP. Informatica Developer, Informatica Analyst, and Data Integration Service are Model repository clients.

The following figure shows how a Model repository client connects to the Model repository database:



1. A Model repository client sends a repository connection request to the master gateway node, which is the entry point to the domain.
2. The Service Manager sends back the host name and port number of the node running the Model Repository Service. In the diagram, the Model Repository Service is running on node A.
3. The repository client establishes a TCP/IP connection with the Model Repository Service process on node A.
4. The Model Repository Service process communicates with the Model repository database over JDBC. The Model Repository Service process stores objects in or retrieves objects from the Model repository database based on requests from the Model repository client.

Note: The Model repository tables have an open architecture. Although you can view the repository tables, never manually edit them through other utilities. Informatica is not responsible for corrupted data that is caused by customer alteration of the repository tables or data within those tables.

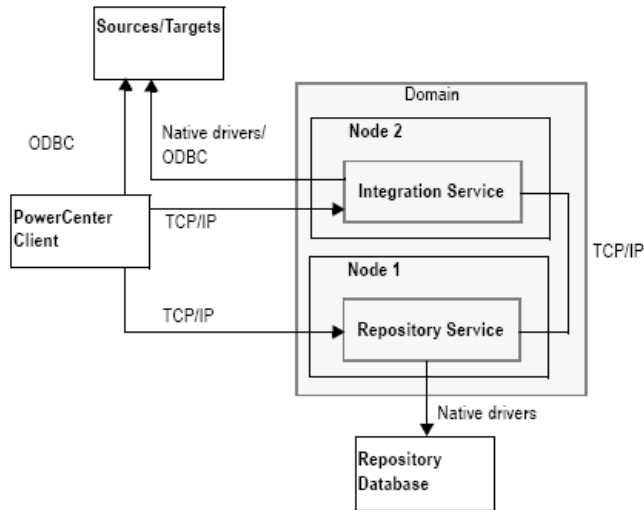
PowerCenter Connectivity

PowerCenter uses the TCP/IP network protocol, native database drivers, ODBC, and JDBC for communication between the following PowerCenter components:

- **PowerCenter Repository Service.** The PowerCenter Repository Service uses native database drivers to communicate with the PowerCenter repository. The PowerCenter Repository Service uses TCP/IP to communicate with other PowerCenter components.
- **PowerCenter Integration Service.** The PowerCenter Integration Service uses native database connectivity and ODBC to connect to source and target databases. The PowerCenter Integration Service uses TCP/IP to communicate with other PowerCenter components.

- **Metadata Manager Service.** Metadata Manager use JDBC and ODBC to access data sources and repositories.
- **PowerCenter Client.** PowerCenter Client uses ODBC to connect to source and target databases. PowerCenter Client uses TCP/IP to communicate with the PowerCenter Repository Service and PowerCenter Integration Service.

The following figure shows an overview of PowerCenter components and connectivity:



The following table lists the drivers used by PowerCenter components:

Component	Database	Driver
PowerCenter Repository Service	PowerCenter Repository	Native
PowerCenter Integration Service	Source Target Stored Procedure Lookup	Native ODBC
Metadata Manager Service	Metadata Manager Repository	JDBC
PowerCenter Client	PowerCenter Repository	Native
PowerCenter Client	Source Target Stored Procedure Lookup	ODBC
Custom Metadata Configurator (Metadata Manager client)	Metadata Manager Repository	JDBC

Repository Service Connectivity

The PowerCenter Repository Service manages the metadata in the PowerCenter repository database. All applications that connect to the repository must connect to the PowerCenter Repository Service. The PowerCenter Repository Service uses native drivers to communicate with the repository database.

The following table describes the connectivity required to connect the Repository Service to the repository and source and target databases:

Repository Service Connection	Connectivity Requirement
PowerCenter Client	TCP/IP
PowerCenter Integration Service	TCP/IP
PowerCenter Repository database	Native database drivers

The PowerCenter Integration Service connects to the Repository Service to retrieve metadata when it runs workflows.

Connecting from PowerCenter Client

To connect to the PowerCenter Repository Service from PowerCenter Client, add a domain and repository in the PowerCenter Client tool. When you connect to the repository from a PowerCenter Client tool, the client tool sends a connection request to the Service Manager on the gateway node. The Service Manager returns the host name and port number of the node where the PowerCenter Repository Service runs. PowerCenter Client uses TCP/IP to connect to the PowerCenter Repository Service.

Connecting to Databases

To set up a connection from the PowerCenter Repository Service to the repository database, configure the database properties in Informatica Administrator. You must install and configure the native database drivers for the repository database on the machine where the PowerCenter Repository Service runs.

Integration Service Connectivity

The PowerCenter Integration Service connects to the repository to read repository objects. The PowerCenter Integration Service connects to the repository through the PowerCenter Repository Service. Use Informatica Administrator to configure an associated repository for the Integration Service.

The following table describes the connectivity required to connect the PowerCenter Integration Service to the platform components, source databases, and target databases:

PowerCenter Integration Service Connection	Connectivity Requirement
PowerCenter Client	TCP/IP
Other PowerCenter Integration Service Processes	TCP/IP

PowerCenter Integration Service Connection	Connectivity Requirement
Repository Service	TCP/IP
Source and target databases	Native database drivers or ODBC Note: The PowerCenter Integration Service on Windows and UNIX can use ODBC drivers to connect to databases. You can use native drivers to improve performance.

The PowerCenter Integration Service includes ODBC libraries that you can use to connect to other ODBC sources. The Informatica installation includes ODBC drivers.

For flat file, XML, or COBOL sources, you can either access data with network connections, such as NFS, or transfer data to the PowerCenter Integration Service node through FTP software. For information about connectivity software for other ODBC sources, refer to your database documentation.

Connecting from the PowerCenter Client

The Workflow Manager communicates with a PowerCenter Integration Service process over a TCP/IP connection. The Workflow Manager communicates with the PowerCenter Integration Service process each time you start a workflow or display workflow details.

Connecting to the PowerCenter Repository Service

When you create a PowerCenter Integration Service, you specify the PowerCenter Repository Service to associate with the PowerCenter Integration Service. When the PowerCenter Integration Service runs a workflow, it uses TCP/IP to connect to the associated PowerCenter Repository Service and retrieve metadata.

Connecting to Databases

Use the Workflow Manager to create connections to databases. You can create connections using native database drivers or ODBC. If you use native drivers, specify the database user name, password, and native connection string for each connection. The PowerCenter Integration Service uses this information to connect to the database when it runs the session.

Note: PowerCenter supports ODBC drivers, such as ISG Navigator, that do not need user names and passwords to connect. To avoid using empty strings or nulls, use the reserved words PmNullUser and PmNullPasswd for the user name and password when you configure a database connection. The PowerCenter Integration Service treats PmNullUser and PmNullPasswd as no user and no password.

PowerCenter Client Connectivity

The PowerCenter Client uses ODBC drivers and native database client connectivity software to communicate with databases. It uses TCP/IP to communicate with the Integration Service and with the repository.

The following table describes the connectivity types required to connect the PowerCenter Client to the Integration Service, repository, and source and target databases:

PowerCenter Client Connection	Connectivity Requirement
Integration Service	TCP/IP
Repository Service	TCP/IP
Databases	ODBC connection for each database

Connecting to the Repository

You can connect to the repository using the PowerCenter Client tools. All PowerCenter Client tools use TCP/IP to connect to the repository through the Repository Service each time you access the repository to perform tasks such as connecting to the repository, creating repository objects, and running object queries.

Connecting to Databases

To connect to databases from the Designer, use the Windows ODBC Data Source Administrator to create a data source for each database you want to access. Select the data source names in the Designer when you perform the following tasks:

- **Import a table or a stored procedure definition from a database.** Use the Source Analyzer or Target Designer to import the table from a database. Use the Transformation Developer, Maplet Designer, or Mapping Designer to import a stored procedure or a table for a Lookup transformation.

To connect to the database, you must also provide your database user name, password, and table or stored procedure owner name.
- **Preview data.** You can select the data source name when you preview data in the Source Analyzer or Target Designer. You must also provide your database user name, password, and table owner name.

Connecting to the Integration Service

The Workflow Manager and Workflow Monitor communicate directly with the Integration Service over TCP/IP each time you perform session and workflow-related tasks, such as running a workflow. When you log in to a repository through the Workflow Manager or Workflow Monitor, the client application lists the Integration Services that are configured for that repository in Informatica Administrator.

Metadata Manager Service Connectivity

To connect to a Metadata Manager repository, the Metadata Manager Service requires a JDBC driver. The Custom Metadata Configurator uses a JDBC driver to connect to the Metadata Manager repository.

JDBC drivers are installed with the Informatica services and the Informatica clients. You can use the installed JDBC drivers to connect to the Metadata Manager repository.

The Informatica installers do not install ODBC drivers or the JDBC-ODBC bridge for the Metadata Manager Service.

Native Connectivity

To establish native connectivity between an application service and a database, you must install the database client software on the machine where the service runs.

The following table describes the syntax for the native connection string for each supported database system:

Database	Connect String Syntax	Example
IBM DB2	<i>dbname</i>	mydatabase
Microsoft SQL Server	<i>servername@dbname</i>	sqlserver@mydatabase
Oracle	<i>dbname.world</i> (same as TNSNAMES entry)	oracle.world
Sybase ASE	<i>servername@dbname</i>	sambrown@mydatabase Note: Sybase ASE servername is the name of the Adaptive Server from the interfaces file.
Teradata	<i>ODBC_data_source_name</i> or <i>ODBC_data_source_name@db_name</i> or <i>ODBC_data_source_name@db_user_name</i>	TeradataODBC TeradataODBC@mydatabase TeradataODBC@sambrown Note: Use Teradata ODBC drivers to connect to source and target databases.

ODBC Connectivity

Open Database Connectivity (ODBC) provides a common way to communicate with different database systems.

To use ODBC connectivity, you must install the following components on the machine hosting the Informatica service or client tool:

- **Database client software.** Install the client software for the database system. This installs the client libraries needed to connect to the database.
Note: Some ODBC drivers contain wire protocols and do not require the database client software.
- **ODBC drivers.** The DataDirect closed 32-bit or 64-bit ODBC drivers are installed when you install the Informatica services. The DataDirect closed 32-bit ODBC drivers are installed when you install the Informatica clients. The database server can also include an ODBC driver.

After you install the necessary components you must configure an ODBC data source for each database that you want to connect to. A data source contains information that you need to locate and access the database, such as database name, user name, and database password. On Windows, you use the ODBC Data Source Administrator to create a data source name. On UNIX, you add data source entries to the `odbc.ini` file found in the system `$ODBCHOME` directory.

When you create an ODBC data source, you must also specify the driver that the ODBC driver manager sends database calls to.

The following table shows the recommended ODBC drivers to use with each database:

Database	ODBC Driver	Requires Database Client Software
Informix	DataDirect Informix Wire Protocol	No
Microsoft Access	Microsoft Access driver	No
Microsoft Excel	Microsoft Excel driver	No
Microsoft SQL Server	DataDirect SQL Server Wire Protocol	No
Netezza	Netezza SQL	Yes
Teradata	Teradata ODBC driver	Yes
SAP HANA	SAP HANA ODBC driver	Yes

JDBC Connectivity

JDBC (Java Database Connectivity) is a Java API that provides connectivity to relational databases. Java-based applications can use JDBC drivers to connect to databases.

The following services and clients use JDBC to connect to databases:

JDBC drivers are installed with the Informatica services and the Informatica clients.

APPENDIX D

Configure the Web Browser

This appendix includes the following topic:

- [Configure the Web Browser, 353](#)

Configure the Web Browser

You can run the Administrator tool in the Microsoft Internet Explorer, Microsoft Edge, Google Chrome, or Safari web browser.

To use the Administrator tool, configure the following options in the browser:

Scripting and ActiveX

Enable the following controls on Microsoft Internet Explorer:

- Active scripting
- Allow programmatic clipboard access
- Run ActiveX controls and plug-ins
- Script ActiveX controls marked safe for scripting

To configure the controls, click **Tools > Internet options > Security > Custom level**.

Trusted sites

If the Informatica domain runs on a network with Kerberos authentication, you must configure the browser to allow access to the Informatica web applications. In Microsoft Internet Explorer, Microsoft Edge, and Google Chrome, add the URL of the Informatica web application to the list of trusted sites. In Safari, add the certificate of the Informatica web application to the keychain. If you are using Chrome version 86.0.42x or later on Windows, you must also set the `AuthServerWhitelist` and `AuthNegotiateDelegateWhitelist` policies.

INDEX

A

- account management
 - overview [60](#)
- accounts
 - changing the password [29](#)
 - managing [28](#)
- activity data
 - Web Services Report [289](#)
- Adabas connections
 - properties [120](#)
- Administrator tool
 - code page [309](#)
 - log errors, viewing [272](#)
 - logs, viewing [267](#)
 - reports [281](#)
- alerts
 - configuring [73](#)
 - description [18](#)
 - managing [73](#)
 - notification email [74](#)
 - subscribing to [74](#)
 - tracking [74](#)
 - viewing [74](#)
- Amazon Redshift connection
 - properties [122](#)
- Amazon S3 connection
 - properties [124](#)
- Analyst Service
 - application service [40](#)
 - custom roles [335](#)
 - log events [274](#)
- application service process
 - disabling [78](#)
 - enabling [78](#)
 - failed state [78](#)
 - port assignment [20](#)
 - standby state [78](#)
 - state [78](#)
 - stopped state [78](#)
- application services
 - Analyst Service [40](#)
 - Content Management Service [40](#)
 - Data Integration Service [40](#)
 - dependencies [66](#)
 - description [20](#)
 - disabling [78](#)
 - enabling [78](#)
 - licenses, assigning [226](#)
 - licenses, unassigning [226](#)
 - Metadata Access Service [40](#)
 - Metadata Manager Service [40](#)
 - Model Repository Service [40](#)
 - overview [40](#)
 - PowerCenter Integration Service [40](#)
 - PowerCenter Repository Service [40](#)

- application services (*continued*)
 - PowerExchange Listener Service [40](#)
 - PowerExchange Logger Service [40](#)
 - removing [79](#)
 - Reporting and Dashboards Service [40](#)
 - Reporting Service [40](#)
 - resilience, configuring [108](#)
 - SAP BW Service [40](#)
 - searching [39](#)
 - Web Services Hub [40](#)
- application sources
 - code page [310](#)
- application targets
 - code page [311](#)
- applications
 - monitoring [242](#)
- ASCII mode
 - overview [303](#)
- audit reports
 - overview [60](#)
- authentication
 - log events [273](#)
- authorization
 - log events [273](#)
 - Service Manager [18](#)
- auto-select
 - network high availability [111](#)
- Average Service Time (property)
 - Web Services Report [289](#)
- Avg DTM Time (property)
 - Web Services Report [289](#)
- Avg. No. of Run Instances (property)
 - Web Services Report [289](#)
- Avg. No. of Service Partitions (property)
 - Web Services Report [289](#)

B

- backing up
 - domain configuration database [81](#)
- backup directory
 - node property [96](#)
- BackupDomain command
 - description [81](#)
- Blaze engine
 - connection properties [140](#)
- blockchain
 - connection properties [126](#)

C

- case study
 - processing ISO 8859-1 data [318](#)
 - processing Unicode UTF-8 data [320](#)

- Cassandra connections
 - properties [128](#)
- catalina.out
 - troubleshooting [264](#)
- category
 - domain log events [273](#)
- changing
 - password for user account [29](#)
- character sizes
 - double byte [307](#)
 - multibyte [307](#)
 - single byte [307](#)
- COBOL
 - connectivity [348](#)
- code page relaxation
 - compatible code pages, selecting [316](#)
 - configuring the Integration Service [316](#)
 - data inconsistencies [315](#)
 - overview [315](#)
 - troubleshooting [316](#)
- code page validation
 - overview [314](#)
 - relaxed validation [315](#)
- code pages
 - Administrator tool [309](#)
 - application sources [310](#)
 - application targets [311](#)
 - choosing [307](#)
 - compatibility diagram [312](#)
 - compatibility overview [307](#)
 - conversion [317](#)
 - Custom transformation [312](#)
 - Data Integration Service process [323](#)
 - descriptions [325](#)
 - domain configuration database [308](#)
 - External Procedure transformation [312](#)
 - flat file sources [310](#)
 - flat file targets [311](#)
 - ID [325](#)
 - lookup database [312](#)
 - Metadata Manager Service [310](#)
 - names [325](#)
 - overview [305](#)
 - pmcmd [309](#)
 - PowerCenter Client [309](#)
 - PowerCenter Integration Service process [309](#), [323](#)
 - relational sources [310](#)
 - relational targets [311](#)
 - relationships [314](#)
 - relaxed validation for sources and targets [315](#)
 - repository [310](#), [323](#)
 - sort order overview [309](#)
 - sources [310](#), [325](#)
 - stored procedure database [312](#)
 - supported code pages [323](#), [325](#)
 - targets [311](#), [325](#)
 - UNIX [306](#)
 - validation [314](#)
 - Windows [306](#)
- command line programs
 - resilience, configuring [110](#)
- compatibility
 - between code pages [307](#)
 - between source and target code pages [316](#)
- compatible
 - defined for code page compatibility [307](#)
- complete history statistics
 - Web Services Report [291](#)
- compute role
 - nodes [93](#)
- Configuration Support Manager
 - using to analyze node diagnostics [298](#)
 - using to review node diagnostics [294](#)
- Confluent Kafka connection
 - Confluent Kafka broker properties [130](#)
 - create using infacmd [131](#)
 - general properties [130](#)
- connect string
 - examples [351](#)
 - syntax [351](#)
- connecting
 - SQL data service [115](#)
- connecting to databases
 - JDBC [350](#)
- Connection
 - details [177](#)
 - properties [177](#)
- connection pooling
 - properties [117](#)
- connection properties
 - Databricks [131](#)
 - blockchain [126](#)
- connection strings
 - native connectivity [351](#)
- connections
 - properties [140](#)
 - adding pass-through security [116](#)
 - creating database connections [113](#)
 - database identifier properties [209](#)
 - deleting [115](#)
 - editing [114](#)
 - Google PubSub [139](#)
 - overview [112](#)
 - pass-through security [115](#)
 - refreshing [113](#)
 - Salesforce Marketing Cloud [193](#)
 - testing [114](#)
 - web services properties [208](#)
- connectivity
 - COBOL [348](#)
 - connect string examples [351](#)
 - Content Management Service [344](#)
 - Data Integration Service [344](#)
 - diagram of [343](#)
 - Informatica Analyst [344](#)
 - Informatica Developer [344](#)
 - Integration Service [348](#)
 - Metadata Manager [350](#)
 - Model Repository Service [344](#)
 - overview [343](#)
 - PowerCenter Client [349](#)
 - PowerCenter Repository Service [348](#)
- Content Management Service
 - application service [40](#)
 - connectivity [344](#)
- Cosmos DB connection
 - creating [178](#)
- CPU detail
 - License Management Report [283](#)
- CPU profile
 - node property [96](#)
- CPU summary
 - License Management Report [282](#)
- CPUs
 - exceeding the limit [282](#)

- creating
 - Cosmos DB connection [178](#)
- custom filters
 - date and time [260](#)
 - elapsed time [260](#)
 - multi-select [260](#)
- custom properties
 - domain [90](#)
- custom roles
 - Analyst Service [335](#)
 - Metadata Manager Service [336](#)
 - Operator [337](#)
 - PowerCenter Repository Service [338](#)

D

- Data Analyzer
 - ODBC (Open Database Connectivity) [343](#)
- Data Integration Service
 - application service [40](#)
 - connectivity [344](#)
 - log events [274](#)
 - recovery [106](#)
- Data Integration Service process
 - supported code pages [323](#)
 - viewing status [95](#)
- Data Integration Services
 - monitoring [237](#)
- data movement mode
 - ASCII [304](#)
 - changing [304](#)
 - description [303](#)
 - effect on session files and caches [304](#)
 - overview [303](#)
 - Unicode [304](#)
- data object caching
 - with pass-through security [116](#)
- database
 - domain configuration [81](#)
- database connections
 - identifier properties [209](#)
 - updating for domain configuration [84](#)
- database drivers
 - Integration Service [343](#)
 - Repository Service [343](#)
- database properties
 - Informatica domain [87](#)
- Databricks connection properties [131](#)
- DataDirect ODBC drivers
 - platform-specific drivers required [351](#)
- deleting
 - connections [115](#)
 - schedules [217](#)
- delimited identifiers
 - database connections [210](#)
- dependencies
 - application services [66](#)
 - grids [66](#)
 - nodes [66](#)
 - viewing for services and nodes [66](#)
- deployed mapping jobs
 - monitoring [243](#)
- detailed statistics
 - monitoring [238](#)
- disable mode
 - PowerCenter Integration Services and Service Processes [78](#)

- domain
 - log event categories [273](#)
 - reports [281](#)
 - user activity, monitoring [281](#)
 - user security [77](#)
- domain configuration
 - description [81](#)
 - log events [273](#)
 - migrating [82](#)
- domain configuration database
 - backing up [81](#)
 - code page [308](#)
 - connection for gateway node [84](#)
 - description [81](#)
 - migrating [82](#)
 - restoring [82](#)
 - secure database [88](#)
 - updating [84](#)
- domain properties
 - Informatica domain [86](#)
- domain reports
 - License Management Report [281](#)
 - running [281](#)
 - Web Services Report [288](#)
- domains
 - multiple [72](#)

E

- editing
 - connections [114](#)
 - schedules [216](#)
- environment variables
 - LANG_C [306](#)
 - LC_ALL [306](#)
 - LC_CTYPE [306](#)
 - NLS_LANG [318](#), [321](#)
 - troubleshooting [79](#)

F

- failover
 - application service [106](#)
 - domain [105](#)
- flat files
 - connectivity [348](#)
 - exporting logs [271](#)
 - source code page [310](#)
 - target code page [311](#)
- folders
 - Administrator tool [75](#)
 - creating [75](#)
 - managing [75](#)
 - objects, moving [76](#)
 - overview [40](#)
 - removing [76](#)
- FTP
 - achieving high availability [111](#)

G

- gateway
 - managing [80](#)
- gateway node
 - configuring [80](#)

- gateway node *(continued)*
 - description [92](#)
 - log directory [80](#)
 - logging [264](#)
- GB18030
 - description [301](#)
- general properties
 - Informatica domain [86](#)
 - license [228](#)
- globalization
 - overview [300](#)
- Google Analytics connections
 - properties [135](#)
- Google BigQuery connection
 - properties [135](#)
- Google BigQuery connections
 - connection modes [137](#)
- Google Cloud Spanner connection
 - properties [137](#)
- Google Cloud Storage connections
 - properties [138](#)
- Google PubSub
 - connection properties [139](#)
- graphics display server
 - requirement [281](#)
- Greenplum connections
 - properties [133](#)
- grids
 - dependencies [66](#)
 - Informatica Administrator tabs [46](#)
 - searching [39](#)
- groups
 - overview [58](#)
- Guaranteed Message Delivery files
 - Log Manager [263](#)

H

- hardware configuration
 - License Management Report [285](#)
- HBase connections
 - MapR-DB properties [149](#)
 - properties [146](#)
- HDFS connections
 - properties [147](#)
- high availability
 - description [25](#), [101](#)
 - failover [105](#)
 - recovery [106](#)
 - restart [105](#)
 - TCP KeepAlive timeout [111](#)
- high availability persistence tables
 - PowerCenter Integration Service [109](#)
- historical statistics
 - monitoring [238](#)
- Hive connections
 - properties [149](#)
- Hive pushdown
 - connection properties [140](#)
- HTTP connections
 - properties [153](#)

I

- IBM DB2
 - connect string syntax [351](#)

- IBM DB2 connections
 - properties [155](#)
- IBM DB2 for i5/OS connections
 - properties [158](#)
- IBM DB2 for z/OS connections
 - properties [161](#)
- identifiers
 - delimited [210](#)
 - regular [210](#)
- IME (Windows Input Method Editor)
 - input locales [303](#)
- IMS connections
 - properties [164](#)
- incremental keys
 - licenses [224](#)
- Informatica Administrator
 - logging in [32](#)
 - Logs tab [56](#)
 - Manage tab [33](#), [38](#)
 - Monitor tab [48](#), [49](#)
 - Navigator [57](#)
 - overview [31](#), [72](#)
 - Reports tab [57](#)
 - searching [57](#)
 - Security page [57](#)
 - service process, enabling and disabling [78](#)
 - Services and Nodes view [39](#)
 - services, enabling and disabling [78](#)
 - tabs, viewing [31](#)
- Informatica Analyst
 - connectivity [344](#)
- Informatica Data Explorer
 - connectivity [344](#)
- Informatica Data Quality
 - connectivity [344](#)
- Informatica Data Services
 - connectivity [344](#)
- Informatica Developer
 - connectivity [344](#)
- Informatica domain
 - alerts [73](#)
 - database properties [87](#)
 - description [17](#)
 - domain configuration database [88](#)
 - domain properties [86](#)
 - general properties [86](#)
 - log and gateway configuration [88](#)
 - multiple domains [72](#)
 - permissions [77](#)
 - privileges [77](#)
 - restarting [85](#)
 - shutting down [85](#)
 - state of operations [106](#)
 - user security [77](#)
- Informatica Network
 - logging in [295](#)
- Information and Content Exchange (ICE)
 - log files [271](#)
- input locales
 - configuring [303](#)
 - IME (Windows Input Method Editor) [303](#)
- Integration Service
 - connectivity [348](#)
 - ODBC (Open Database Connectivity) [343](#)

J

- JD Edwards EnterpriseOne connection
 - properties [171](#)
- JDBC (Java Database Connectivity)
 - overview [352](#)
- JDBC connections
 - properties [166](#)
- JDBC drivers
 - Data Analyzer [343](#)
 - installed drivers [350](#)
 - Metadata Manager [343](#)
 - Metadata Manager connection to databases [350](#)
 - PowerCenter domain [343](#)
 - Reference Table Manager [343](#)
- JDBC V2 connection
 - properties [169](#)
- job scheduling
 - overview [213](#)
- job status
 - domain failover [259](#)
- jobs
 - monitoring [238](#)

K

- Kafka connection
 - create using infacmd [175](#)
 - general properties [173](#)
 - Kafka broker properties [173](#)
- Kerberos authentication
 - troubleshooting [33](#)
- Kudu connection
 - properties [175](#)

L

- LANG_C environment variable
 - setting locale in UNIX [306](#)
- LC_ALL environment variable
 - setting locale in UNIX [306](#)
- LDAP connection
 - properties [176](#)
- license
 - assigning to a service [226](#)
 - creating [225](#)
 - details, viewing [228](#)
 - general properties [228](#)
 - Informatica Administrator tabs [46](#)
 - keys [224](#)
 - license file [225](#)
 - log events [273](#), [276](#)
 - managing [223](#)
 - removing [227](#)
 - unassigning from a service [226](#)
 - updating [227](#)
 - validation [223](#)
- license keys
 - incremental [224](#), [227](#)
 - original [224](#)
- License Management Report
 - CPU detail [283](#)
 - CPU summary [282](#)
 - emailing [287](#)
 - hardware configuration [285](#)
 - licensed options [286](#)

- License Management Report (*continued*)
 - licensing [282](#)
 - multibyte characters [287](#)
 - node configuration [285](#)
 - repository summary [284](#)
 - running [281](#), [286](#)
 - Unicode font [287](#)
 - user detail [284](#)
 - user summary [284](#)
- license usage
 - log events [273](#)
- licensed options
 - License Management Report [286](#)
- licensing
 - License Management Report [282](#)
 - log events [275](#)
 - managing [223](#)
- licensing logs
 - log events [223](#)
- linked domain
 - multiple domains [72](#)
- Listener Service
 - log events [274](#)
- locales
 - overview [302](#)
- localhost.txt
 - troubleshooting [264](#)
- Log Agent
 - description [262](#)
 - log events [273](#)
- log and gateway configuration
 - Informatica domain [88](#)
- log directory
 - for gateway node [80](#)
 - location, configuring [265](#)
- log errors
 - Administrator tool [272](#)
- log event files
 - description [263](#)
 - purging [266](#)
- log events
 - authentication [273](#)
 - authorization [273](#)
 - code [273](#)
 - components [273](#)
 - description [263](#)
 - details, viewing [268](#)
 - domain [273](#)
 - domain configuration [273](#)
 - domain function categories [273](#)
 - exporting with Mozilla Firefox [270](#)
 - licensing [273](#), [275](#), [276](#)
 - licensing logs [223](#)
 - licensing usage [273](#)
 - Log Agent [273](#)
 - Log Manager [273](#)
 - message [273](#)
 - message code [273](#)
 - node [273](#)
 - node configuration [273](#)
 - PowerCenter Repository Service [276](#)
 - saving [270](#)
 - security audit trail [276](#)
 - Service Manager [273](#)
 - service name [273](#)
 - severity levels [273](#)
 - thread [273](#)
 - time zone [267](#)

- log events (*continued*)
 - timestamps [273](#)
 - user activity [277](#)
 - user management [273](#)
 - viewing [267](#)
 - Web Services Hub [277](#)
 - workflow [257](#)
- log files
 - Mapping tasks [279](#)
- Log Manager
 - architecture [263](#)
 - catalina.out [264](#)
 - configuring [267](#)
 - directory location, configuring [265](#)
 - domain log events [273](#)
 - log event components [273](#)
 - log events [273](#)
 - log events, purging [266](#)
 - log events, saving [270](#)
 - logs, viewing [267](#)
 - message [273](#)
 - message code [273](#)
 - node [273](#)
 - node.log [264](#)
 - PowerCenter Integration Service log events [275](#)
 - PowerCenter Repository Service log events [276](#)
 - ProcessID [273](#)
 - purge properties [266](#)
 - recovery [264](#)
 - SAP NetWeaver BI log events [276](#)
 - security audit trail [276](#)
 - service name [273](#)
 - severity levels [273](#)
 - thread [273](#)
 - time zone [267](#)
 - timestamp [273](#)
 - troubleshooting [264](#)
 - user activity log events [277](#)
 - using [262](#)
- Logger Service
 - log events [275](#)
- logical CPUs
 - calculation [282](#)
- logical data objects
 - monitoring [246](#)
- login
 - troubleshooting [33](#)
- logs
 - components [273](#)
 - configuring [265](#)
 - domain [273](#)
 - location [265](#)
 - PowerCenter Integration Service [275](#)
 - PowerCenter Repository Service [276](#)
 - purging [266](#)
 - SAP BW Service [276](#)
 - saving [270](#)
 - user activity [277](#)
 - viewing [267](#)
 - workflow [257](#)
- Logs tab
 - Informatica Administrator [56](#)
- lookup databases
 - code pages [312](#)

M

- Manage tab
 - Connections view [46](#)
 - Informatica Administrator [33, 38](#)
 - Navigator [33, 38](#)
 - Schedules view [47](#)
 - Services and Nodes view [38](#)
- managing
 - accounts [28](#)
 - user accounts [28](#)
- Mapping task
 - log files [279](#)
- master gateway node
 - description [92](#)
- Maximum CPU Run Queue Length
 - node property [96](#)
- Maximum Memory Percent
 - node property [96](#)
- Maximum Processes
 - node property [96](#)
- Maximum Restart Attempts (property)
 - Informatica domain [79](#)
- message code
 - Log Manager [273](#)
- Messaging connection
 - Confluent Kafka connection [129](#)
 - Kafka connection [172](#)
- metadata
 - adding to repository [317](#)
 - choosing characters [317](#)
- Metadata Access Service
 - application service [40](#)
- Metadata Manager
 - connectivity [350](#)
 - ODBC (Open Database Connectivity) [343](#)
- Metadata Manager Service
 - application service [40](#)
 - code page [310](#)
 - custom roles [336](#)
 - log events [275](#)
- Microsoft Azure Data Lake Storage Gen1 connection
 - properties [179](#)
- Microsoft Azure Data Lake Storage Gen2 connection
 - properties [180](#)
- Microsoft Azure SQL Data Warehouse connection
 - properties [181](#)
- Microsoft SQL Server
 - connect string syntax [351](#)
- migrate
 - domain configuration [82](#)
- Model Repository Service
 - application service [40](#)
 - connectivity [344](#)
 - log events [275](#)
- Monitor tab
 - Informatica Administrator [48, 49](#)
- monitoring
 - applications [242](#)
 - configuring [233](#)
 - Data Integration Services [237](#)
 - deployed mapping jobs [243](#)
 - description [231](#)
 - detailed statistics [238](#)
 - exporting summary statistics [236](#)
 - historical statistics [238](#)
 - jobs [238](#)
 - logical data objects [246](#)

- monitoring (*continued*)
 - preferences, configuring [234](#)
 - reports [53](#)
 - setup [232](#)
 - SQL data services [247](#)
 - statistics [52](#)
 - summary statistics [235](#), [238](#)
 - viewing summary statistics [236](#)
 - web services [250](#)
 - workflows [251](#)
- monitoring Model Repository Service
 - application service [40](#)
- MS SQL Server connections
 - properties [182](#)
- multibyte data
 - entering in PowerCenter Client [303](#)

N

- Navigator
 - Manage tab [33](#), [38](#)
 - searching [39](#)
 - Security page [57](#)
- Netezza connections
 - properties [186](#)
- network
 - high availability [111](#)
- NLS_LANG
 - setting locale [318](#), [321](#)
- node configuration
 - License Management Report [285](#)
 - log events [273](#)
- node configuration file
 - location [95](#)
- node diagnostics
 - analyzing [298](#)
 - downloading [297](#)
- node properties
 - backup directory [96](#)
 - configuring [96](#)
 - CPU Profile [96](#)
 - maximum CPU run queue length [96](#)
 - maximum memory percent [96](#)
 - maximum processes [96](#)
- node roles
 - compute [93](#)
 - service [93](#)
 - updating [94](#)
- node.log
 - troubleshooting [264](#)
- nodemeta.xml
 - for gateway node [80](#)
 - location [95](#)
- nodes
 - adding to Informatica Administrator [95](#)
 - configuring [96](#)
 - defining [95](#)
 - dependencies [66](#)
 - description [17](#), [92](#)
 - gateway [80](#), [92](#)
 - host name and port number, removing [96](#)
 - Informatica Administrator tabs [45](#)
 - Log Manager [273](#)
 - port number [96](#)
 - removing [100](#)
 - restarting [98](#)
 - roles [93](#)

- nodes (*continued*)
 - searching [39](#)
 - shutting down [98](#)
 - starting [98](#)
 - TCP/IP network protocol [343](#)
 - types [92](#)
 - worker [92](#)

O

- OData connections
 - properties [187](#)
- ODBC (Open Database Connectivity)
 - DataDirect driver issues [351](#)
 - establishing connectivity [351](#)
 - Integration Service [343](#)
 - Metadata Manager [343](#)
 - PowerCenter Client [343](#)
 - requirement for PowerCenter Client [349](#)
- ODBC connections
 - properties [188](#)
- operating mode
 - effect on resilience [110](#)
- operating system profiles
 - overview [59](#)
- Operator}
 - custom roles [337](#)
- Oracle
 - connect string syntax [351](#)
 - setting locale with NLS_LANG [318](#), [321](#)
- Oracle connections
 - properties [189](#)
- original keys
 - licenses [224](#)
- overview
 - connections [112](#)

P

- pass-through security
 - adding to connections [116](#)
 - connecting to SQL data service [115](#)
 - enabling caching [116](#)
 - web service operation mappings [115](#)
- password
 - changing for a user account [29](#)
- Percent Partitions in Use (property)
 - Web Services Report [289](#)
- pmcmd
 - code page issues [309](#)
 - communicating with PowerCenter Integration Service [309](#)
- PmNullPasswd
 - reserved word [349](#)
- PmNullUser
 - reserved word [349](#)
- port
 - application service [20](#)
 - node [96](#)
 - node maximum [96](#)
 - node minimum [96](#)
 - range for service processes [96](#)
- PowerCenter
 - connectivity [343](#)
- PowerCenter Client
 - code page [309](#)
 - connectivity [349](#)

- PowerCenter Client (*continued*)
 - multibyte characters, entering [303](#)
 - ODBC (Open Database Connectivity) [343](#)
 - resilience [102](#)
 - TCP/IP network protocol [343](#)
- PowerCenter domains
 - connectivity [346](#)
 - TCP/IP network protocol [343](#)
- PowerCenter Integration Service
 - application service [40](#)
 - enabling and disabling [78](#)
 - failover configuration [109](#)
 - high availability persistence tables [109](#)
 - log events [275](#)
 - recovery [106](#)
 - recovery configuration [109](#)
 - resilience [103](#)
 - state of operations [106](#)
- PowerCenter Integration Service process
 - code page [309](#)
 - enabling and disabling [78](#)
 - restart, configuring [79](#)
 - supported code pages [323](#)
 - viewing status [95](#)
- PowerCenter Repository Service
 - application service [40](#)
 - connectivity requirements [348](#)
 - custom roles [338](#)
 - log events [276](#)
 - recovery [106](#)
 - resilience [103](#)
 - state of operations [106](#)
- PowerCenter security
 - managing [57](#)
- PowerExchange Listener Service
 - application service [40](#)
- PowerExchange Logger Service
 - application service [40](#)
- preferences
 - monitoring [234](#)
- process identification number
 - Log Manager [273](#)
- ProcessID
 - Log Manager [273](#)
 - message code [273](#)
- purge properties
 - Log Manager [266](#)

R

- recovery
 - Data Integration Service [106](#)
 - high availability [106](#)
 - Integration Service [106](#)
 - PowerCenter Repository Service [106](#)
- regular identifiers
 - database connections [210](#)
- Reporting and Dashboards Service
 - application service [40](#)
- Reporting Service
 - application service [40](#)
- reports
 - Administrator tool [281](#)
 - domain [281](#)
 - License [281](#)
 - monitoring [53](#)
 - Web Services [281](#)

- Reports tab
 - Informatica Administrator [57](#)
- repositories
 - backup directory [96](#)
 - code pages [310](#)
 - supported code pages [323](#)
 - Unicode [301](#)
 - UTF-8 [301](#)
- repository metadata
 - choosing characters [317](#)
- repository summary
 - License Management Report [284](#)
- resilience
 - application service [103](#)
 - application service configuration [108](#)
 - command line program configuration [110](#)
 - in exclusive mode [110](#)
 - PowerCenter Client [102](#)
 - PowerCenter Integration Service [103](#)
 - PowerCenter Repository Service [103](#)
 - TCP KeepAlive timeout [111](#)
- Resource Manager Service
 - log events [276](#)
 - system services [44](#)
- resource provision thresholds
 - setting for nodes [96](#)
- restart
 - application service [106](#)
 - configuring for PowerCenter Integration Service processes [79](#)
- restoring
 - domain configuration database [82](#)
- roles
 - nodes [93](#)
 - overview [59](#)
- run-time statistics
 - Web Services Report [290](#)

S

- Salesforce Marketing Cloud
 - connection properties [193](#)
- SAP BW Service
 - application service [40](#)
 - log events [276](#)
- SAP connections
 - properties [194](#)
- Scheduler Service
 - log events [277](#)
- schedules
 - creating schedules [214](#)
 - deleting [217](#)
 - editing [216](#)
 - overview [213](#)
- Search section
 - Informatica Administrator [57](#)
- security
 - audit trail, viewing [276](#)
 - permissions [77](#)
 - privileges [77](#)
- Security page
 - Informatica Administrator [57](#)
 - Navigator [57](#)
- Sequential connections
 - properties [197](#)
- Service Manager
 - authorization [18](#)
 - description [18](#)

- Service Manager (*continued*)
 - log events [273](#)
- service name
 - log events [273](#)
- service role
 - nodes [93](#)
- services
 - searching [39](#)
- services and nodes
 - viewing dependencies [66](#)
- Services and Nodes view
 - Informatica Administrator [39](#)
- sessions
 - sort order [309](#)
- severity
 - log events [273](#)
- Show Custom Properties (property)
 - user preference [29](#)
- shutting down
 - Informatica domain [85](#)
- SMTP configuration
 - alerts [73](#)
- Snowflake connection
 - properties [199](#)
- sort order
 - code page [309](#)
- source databases
 - code page [310](#)
- sources
 - code pages [310](#), [325](#)
- Spark deploy mode
 - Hadoop connection properties [140](#)
- Spark engine
 - connection properties [140](#)
- Spark Event Log directory
 - Hadoop connection properties [140](#)
- Spark execution parameters
 - Hadoop connection properties [140](#)
- Spark HDFS staging directory
 - Hadoop connection properties [140](#)
- SQL data services
 - monitoring [247](#)
- stack traces
 - viewing [268](#)
- state of operations
 - domain [106](#)
 - PowerCenter Integration Service [106](#)
 - PowerCenter Repository Service [106](#)
- statistics
 - for monitoring [52](#)
 - Web Services Hub [288](#)
- stopping
 - Informatica domain [85](#)
- stored procedures
 - code pages [312](#)
- Subscribe for Alerts
 - user preference [29](#)
- subset
 - defined for code page compatibility [307](#)
- summary statistics
 - monitoring [235](#)
- superset
 - defined for code page compatibility [307](#)
- Sybase ASE
 - connect string syntax [351](#)
- system locales
 - description [302](#)

- system services
 - Resource Manager Service [44](#)

T

- Tableau V3 connection
 - properties [203](#)
- target databases
 - code page [311](#)
- targets
 - code pages [311](#), [325](#)
- tasks
 - states [254](#)
- TCP KeepAlive timeout
 - high availability [111](#)
- TCP/IP network protocol
 - nodes [343](#)
 - PowerCenter Client [343](#)
 - PowerCenter domains [343](#)
 - requirement for Integration Service [349](#)
- Teradata
 - connect string syntax [351](#)
- Teradata Parallel Transporter connections
 - properties [200](#)
- testing
 - database connections [114](#)
- thread identification
 - Logs tab [273](#)
- threads
 - Log Manager [273](#)
- time zone
 - Log Manager [267](#)
- timestamps
 - Log Manager [273](#)
- troubleshooting
 - catalina.out [264](#)
 - code page relaxation [316](#)
 - environment variables [79](#)
 - Kerberos authentication [33](#)
 - localhost.txt [264](#)
 - logging in [33](#)
 - node.log [264](#)
- Twitter Streaming connections
 - properties [204](#)

U

- UCS-2
 - description [301](#)
- Unicode
 - GB18030 [301](#)
 - repositories [301](#)
 - UCS-2 [301](#)
 - UTF-16 [301](#)
 - UTF-32 [301](#)
 - UTF-8 [301](#)
- Unicode mode
 - overview [303](#)
- UNIX
 - code pages [306](#)
- UNIX environment variables
 - LANG_C [306](#)
 - LC_ALL [306](#)
 - LC_CTYPE [306](#)
- user accounts
 - changing the password [29](#)

- user accounts (*continued*)
 - managing [28](#)
- user activity
 - log event categories [277](#)
- user detail
 - License Management Report [284](#)
- user locales
 - description [303](#)
- user management
 - log events [273](#)
- user preferences
 - description [29](#)
- user summary
 - License Management Report [284](#)
- users
 - license activity, monitoring [281](#)
 - overview [58](#)
- UTF-16
 - description [301](#)
- UTF-32
 - description [301](#)
- UTF-8
 - description [301](#)
 - repository [310](#)

V

- validating
 - code pages [314](#)
 - licenses [223](#)
- viewing
 - dependencies for services and nodes [66](#)
- VSAM connections
 - properties [205](#)

W

- web connections
 - properties [153](#)

- web services
 - monitoring [250](#)
- Web Services Hub
 - application service [25, 40](#)
 - log events [277](#)
 - statistics [288](#)
- Web Services Report
 - activity data [289](#)
 - Average Service Time (property) [289](#)
 - Avg DTM Time (property) [289](#)
 - Avg. No. of Run Instances (property) [289](#)
 - Avg. No. of Service Partitions (property) [289](#)
 - complete history statistics [291](#)
 - contents [289](#)
 - Percent Partitions in Use (property) [289](#)
 - run-time statistics [290](#)
- Within Restart Period (property)
 - Informatica domain [79](#)
- worker node
 - configuring as gateway [80](#)
 - description [92](#)
- workflow recovery
 - overview [256](#)
 - running [257](#)
- workflows
 - aborting [256](#)
 - canceling [256](#)
 - logs [257](#)
 - monitoring [251](#)
 - recovering [257](#)
 - states [253](#)

X

- X Virtual Frame Buffer
 - for License Report [281](#)
 - for Web Services Report [281](#)
- XML
 - exporting logs in [271](#)