Informatica® Intelligent Cloud Services
Fall 2020 January

Administrator

Informatica Intelligent Cloud Services Administrator
Fall 2020 January
January 2021

Publication Date: 2021-01-11

Table of Contents

Preface

Use *Administrator* to learn how to set up and maintain your Informatica Intelligent Cloud Services℠ organization and sub-organizations. Learn how to manage licenses, configure ecosystem and SAML single sign-on, configure source control, manage users, configure object permissions, configure runtime environments and Secure Agent services, create schedules, manage bundles, monitor events, configure elastic clusters, and configure file servers.

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 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 Intelligent Cloud Services web site

You can access the Informatica Intelligent Cloud Services web site at http://www.informatica.com/cloud. This site contains information about Informatica Cloud integration services.

Informatica Intelligent Cloud Services Communities

Use the Informatica Intelligent Cloud Services Community to discuss and resolve technical issues. You can also find technical tips, documentation updates, and answers to frequently asked questions.

Access the Informatica Intelligent Cloud Services Community at:

https://network.informatica.com/community/informatica-network/products/cloud-integration

Developers can learn more and share tips at the Cloud Developer community:

https://network.informatica.com/community/informatica-network/products/cloud-integration/cloud-developers

Informatica Intelligent Cloud Services Marketplace

Visit the Informatica Marketplace to try and buy Data Integration Connectors, templates, and mapplets:

https://marketplace.informatica.com/

Data Integration connector documentation

You can access documentation for Data Integration Connectors at the Documentation Portal. To explore the Documentation Portal, visit https://docs.informatica.com.

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 Intelligent Cloud Services Trust Center

The Informatica Intelligent Cloud Services Trust Center provides information about Informatica security policies and real-time system availability.

You can access the trust center at https://www.informatica.com/trust-center.html.

Subscribe to the Informatica Intelligent Cloud Services Trust Center to receive upgrade, maintenance, and incident notifications. The Informatica Intelligent Cloud Services Status page displays the production status of all the Informatica cloud products. All maintenance updates are posted to this page, and during an outage, it will have the most current information. To ensure you are notified of updates and outages, you can subscribe to receive updates for a single component or all Informatica Intelligent Cloud Services components. Subscribing to all components is the best way to be certain you never miss an update.

To subscribe, go to https://status.informatica.com/ and click **SUBSCRIBE TO UPDATES**. You can then choose to receive notifications sent as emails, SMS text messages, webhooks, RSS feeds, or any combination of the four.

Informatica Global Customer Support

You can contact a Customer Support Center by telephone or online.

For online support, click **Submit Support Request** in Informatica Intelligent Cloud Services. You can also use Online Support to log a case. Online Support requires a login. You can request a login at https://network.informatica.com/welcome.

The telephone numbers for Informatica Global Customer Support are available from the Informatica web site at https://www.informatica.com/services-and-training/support-services/contact-us.html.

CHAPTER 1

Introducing Administrator

Administrator provides organization management capabilities across Informatica Intelligent Cloud Services.

Use Administrator to manage the following aspects of your organization:

Organization and sub-organizations

Configure settings for your organization and sub-organizations such as password requirements, trusted IP addresses, connection properties storage, time zone and email notification settings for Data Integration tasks, CLAIRE™ recommendation preferences, and Enterprise Data Catalog settings. Create and manage sub-organizations.

For information about organizations and sub-organizations, see Chapter 2, “Organizations” on page 14.

Licenses

View your organization's licenses, manage sub-organization licenses, and view metering information which includes job limits and usage information.

For information about licenses and metering, see Chapter 3, “Licenses” on page 25.

Ecosystem and SAML single sign-on

Configure single-sign on settings for Microsoft Azure. Enable single sign-on capability for a SAML third-party identity provider.

For information about Microsoft Azure single sign-on settings, see Chapter 4, “Ecosystem single sign-on” on page 28. For information about enabling and configuring SAML single sign-on, see Chapter 5, “SAML single sign-on” on page 30.

Source control and Secure Agent service upgrade settings

Configure source control to enable version management for projects, folders, and assets.

Configure upgrade error handling and upgrade restart schedules for some Secure Agent services.

For more information about source control and Secure Agent service upgrade settings, see Chapter 7, “Source control and service upgrade settings” on page 43.

Users, user groups, and user roles

Create and configure individual user accounts to allow access to your organization. Create groups of users that can perform the same tasks. Create and configure roles to define the privileges for your users and user groups.

For information about users and user groups, see Chapter 8, “Users and user groups” on page 51. For information about user roles, see Chapter 9, “User roles” on page 64.

Permissions

Configure the access rights that users and user groups have for objects such as Secure Agents, Secure Agent groups, connections, and schedules.

For information about permissions and configuring permissions, see Chapter 10, “Permissions” on page 79.

Runtime environments

Download and install Secure Agents. Create and configure Secure Agent groups.

For information about Secure Agents and Secure Agent groups, see Chapter 11, “Runtime environments” on page 83. For information about downloading and installing a Secure Agent, see Chapter 14, “Secure Agent installation” on page 146.

Serverless runtime environments

Use a runtime environment that Data Integration manages to reduce maintenance overhead. The serverless runtime environment can run mapping tasks that are based on a mapping or an elastic mapping.

Note: To use a serverless runtime environment, you must have a private cloud on the AWS cloud platform.

For information about serverless runtime environments, see Chapter 12, “Serverless runtime environments” on page 104.

Secure Agent services

Configure settings for the microservices that the Secure Agent uses for data processing such as the Elastic Server, CIH Processor, Data Integration Server, EDC Search Agent, and Process Server.

For information about Secure Agent services and their configuration, see Chapter 13, “Secure Agent services” on page 114.

Elastic clusters

Manage the ephemeral clusters that your organization can use to process data integration jobs on the Serverless Spark engine.

For information about elastic clusters, see *Data Integration Elastic Administration*.

Schedules

Create schedules to run tasks or taskflows at specified times or at regular intervals. Define a blackout period in which no scheduled tasks or jobs in your organization can run.

For information about schedules and organization blackout periods, see Chapter 15, “Schedules” on page 153.

Add-on bundles

Install, copy, upgrade, and uninstall sets of related mappings, mapping tasks, mapplets, and Visio templates that Data Integration users can use in data integration projects.

For information about managing add-on bundles, see Chapter 16, “Bundle management” on page 158.

Event monitoring

Monitor events for the assets, licenses, users, and Secure Agents in your organization through the asset and security logs.

For information about asset and security logs, see Chapter 17, “Event monitoring” on page 161.

File transfer

Configure your organization's file server to securely send and receive files from a business partner's remote server. Configure a connection, and then send the files to your partners using the Informatica Intelligent Cloud Services REST API.

For information about file servers and file transfer, see Chapter 18, “File transfer” on page 163.

Editing your user profile

Your user profile contains the details of your Informatica Intelligent Cloud Services user account.

You can update the following information in your profile:

- Email address
- Time zone (used in the job execution time stamps on the **All Jobs**, **Running Jobs**, **My Jobs**, **Import/Export Logs**, and **My Import/Export Logs** pages)
- Password
- Security question

To edit your user profile:

1. Click the **User** icon in the top right corner of the Informatica Intelligent Cloud Services window and then select **Profile**.
2. On the **Profile** page, add or edit personal information such as your name, job title, phone number, email address, and time zone.
3. Optionally, change your password or security question.
4. Click **Save**.

CHAPTER 2

Organizations

An organization is a secure area within the Informatica Intelligent Cloud Services repository that stores your licenses, user accounts, data integration assets such as mappings and tasks, and information about jobs and security. Based on your license, you might have access to one organization or to a parent organization and one or more sub-organizations.

The administrator of an organization maintains the organization and sub-organizations.

Log in to Informatica Intelligent Cloud Services as an administrator to set up your organization, create and manage schedules, and monitor activities related to assets and security.

Setting up an organization

When you set up an organization, you configure the organization properties, sub-organizations, licenses, runtime environments, and user accounts.

To set up your company's organization, perform the following steps:

1. Configure organization properties such as the organization name and address, authentication information, and notification email addresses.
2. Optionally, create one or more sub-organizations.
3. Verify that your organization has the appropriate licenses, and configure licenses for your sub-organizations.
4. Configure runtime environments and Secure Agents.
5. Set up users, user groups, and roles.

You might also need to download and install non-native connectors for your organization. For example, if users in your organization create tasks that read data from Teradata tables, you need to download and install the add-on connector for Teradata. For more information about downloading and installing add-on connectors, see *Connections*.

Organization properties

Configure properties for your organization or sub-organizations on the Organization page. To access the Organization page, in Administrator, select **Organization**.

The following image shows the Organization page:

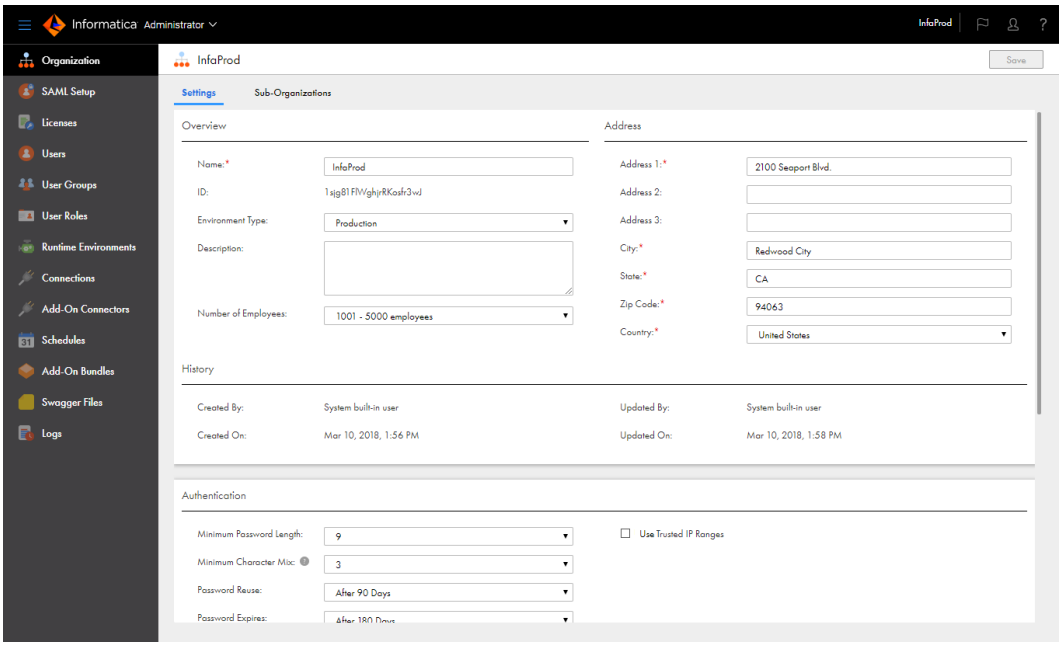

You can configure the following properties:

- General properties such as organization name, description, number of employees, and address information.
- Authentication information and connection properties storage.
- Data Integration service properties such as the time zone and default addresses for email notifications.
- CLAIRE™ recommendation preferences. If enabled, CLAIRE provides design time recommendations based on collected metadata.
- Enterprise Data Catalog integration properties such as the URL of the Enterprise Data Catalog Service, runtime environment that reads data from Enterprise Data Catalog, and Enterprise Data Catalog user name and password.

Organization general properties

You can configure general properties for your organization and sub-organizations. General properties include information such as the organization name, ID, description, address, and number of employees. History information for the organization is also displayed in the general properties.

The general properties include the following information:

Overview information

The following table describes the overview properties:

Property	Description
Name	Name of the organization. If you change the organization name, the new name appears on the **Organization** menu after you log out and log back in.
ID	ID assigned to your organization when it was created. You cannot change an organization ID.
Parent organization ID	When you view a sub-organization, this property displays the ID assigned to the parent organization. You cannot change an organization ID.
Environment type	Environment type, either Development, Production, QA, or Sandbox.
Description	Optional description of the organization.
Number of employees	Number of employees in the organization.
Deny parent organization access to this sub-organization	When this option is checked, users in the parent organization cannot switch from the parent organization to the sub-organization. Users in the parent organization with the appropriate privileges can make only the following changes to the sub-organization: - Enable and disable the sub-organization - Update the sub-organization licenses - Edit the sub-organization properties such as the organization description and CLAIRE recommendation preferences This option is displayed on the **Organization** page for sub-organizations. This option can be changed when an administrator in the sub-organization logs in to the sub-organization. This option is read-only when a parent organization administrator views the organization properties for the sub-organization. This option is unchecked by default.

Address information

Use the address properties to specify the street address of the organization.

History information

The organization history information displays the date and time that the organization was created, the user who created the organization, the date and time that the organization was last updated, and the user who last updated the organization. Informatica Intelligent Cloud Services updates the history information when you make changes to the organization.

Authentication properties

You can configure authentication properties for your organization and sub-organizations. Authentication properties control password restrictions and IP address filtering.

Password restrictions are enforced when users create or change their passwords. If you change the password expiration date from "never" to a number of days, then users with passwords that are older than the number of days will be required to change their passwords the next time that they log in to Informatica Intelligent Cloud Services.

The following table describes the authentication properties:

Property	Description
Minimum Password Length	Minimum password length required for a valid password. Must be a number between 4 and 12 characters.
Minimum Character Mix	Minimum number of character types required for a valid password. Passwords can contain a mix of the following character sets: - Lowercase alphabetic characters - Uppercase alphabetic characters - Numeric characters - Special characters For example, if you set **Minimum Character Mix** to 1, then passwords must contain at least one of the character sets. If you set **Minimum Character Mix** to 2, then passwords must contain at least two of the character sets.
Password Reuse	Controls whether users can reuse passwords.
Password Expires	Determines how often users must reset their passwords.
Session Idle Timeout	Amount of time before a user's session times out due to inactivity. Informatica Intelligent Cloud Services displays a warning message to the user 60 seconds before the user is logged out. Default is 30 minutes.
Use Trusted IP Ranges	Enables IP address filtering. IP address filtering uses trusted IP address ranges in addition to account passwords to prevent unauthorized users from accessing your organization. When you enable IP address filtering, a user with a valid login must also have an IP address within the range of trusted IP addresses, or the user cannot log in to your organization. When you enable this option, you must also enter one or more trusted IP address ranges.
Allowed Trusted IP Ranges	The trusted ranges of IP addresses from which users can log in to access the organization. Informatica Intelligent Cloud Services supports IP address formats in IP version 4 (IPv4) and version 6 (IPv6). Fields for the trusted IP address range appear when you enable IP address filtering. To enter additional address ranges, click **+**. **Note:** If you enter an invalid IP address range, users cannot access your organization. Contact your network administrator for valid IP address ranges.

Connection properties storage

You can configure where to store the connection properties for your organization and sub-organizations. To specify where to store the connection properties, configure the **Connection Credentials** on the **Organization** page.

You can store connection properties in either of the following locations:

Informatica Cloud

When you store connection properties with Informatica Intelligent Cloud Services, the connection properties are always available. Informatica Intelligent Cloud Services backs up connection properties regularly as part of standard backup procedures.

Local Secure Agent

You might store connection properties with a local Secure Agent if you need the connection properties to reside within your firewall. To store connection properties with a local Secure Agent, the organization must have the Externalize Connections license. When you enable this option, the properties for all connections that are listed on the **Connections** page are stored with the local agent.

When you store properties with a local Secure Agent, the Secure Agent must be running so that tasks can run and users can work with connections. Back up connection properties regularly to prevent loss of data. A best practice is to back up connection properties after you change the location or the encryption key for connection properties.

Connection properties are stored in the following directory:

```
<Secure Agent installation directory>/apps/Data_Integration_Server/data
```

Informatica Intelligent Cloud Services generates an encryption key to secure connection properties stored with a Secure Agent. You can use a randomly generated password or you can enter a custom password as the basis for the encryption key.

Use a custom password when you want to update the encryption key periodically. You can change the custom password when you want to update the encryption key.

You can change where you want to store connection properties. When you do this, Informatica Intelligent Cloud Services moves the connection properties to the appropriate location. For example, your license expires, so you configure the organization to store connections on the cloud. Informatica Intelligent Cloud Services moves the connection properties from the local Secure Agent to Informatica Intelligent Cloud Services.

Data Integration service properties

Data Integration service properties are used by Data Integration. Configure these properties to set the time zone and default email addresses for job notifications.

You can set the following Data Integration service properties:

Jobs properties

The following table describes the jobs properties:

Property	Description
Schedule Offset	A small amount of time that is added to schedule start times to help prevent server overload at standard schedule start times. An organization has a single schedule offset that is applied to all schedules. The schedule offset does not affect the start time of manually started tasks or taskflows. You cannot change the schedule offset. Even though it is not displayed in the schedule details, the schedule offset for your organization is added to the time range configured for all schedules. This ensures that scheduled tasks run as often as expected. For example, you configure a schedule to run every hour from 8:00 a.m. to 12:00 p.m., and the schedule offset for your organization is 15 seconds. Your schedule runs at 8:00:15, 9:00:15, 10:00:15, 11:00:15, and 12:00:15.
Time Zone	Time zone used to display job execution time stamps in email notifications.

Email notification properties

Configure the email notification properties to set the default email addresses to use for job failure, warning, and success messages. Enter one or more valid email addresses. Separate email addresses with a comma (,) or semicolon (;).

You can also set email notification properties at the task level. When you set email notifications in a task or taskflow, Informatica Intelligent Cloud Services sends email to the addresses in the task or taskflow instead of the addresses configured for the organization.

CLAIRE recommendation preferences

Enable CLAIRE recommendations to allow in-product recommendations for mapping design based on analysis of metadata from your organization's assets and assets from other Informatica Intelligent Cloud Services organizations. The metadata collected and processed by the CLAIRE engine is anonymous.

The default setting for CLAIRE recommendations is "Enabled." When you disable CLAIRE recommendations, recommendations are disabled for all users within your organization. You can enable or disable recommendations for your organization at any time.

Enable and disable CLAIRE recommendations for suborganizations from within the suborganization.

When you enable CLAIRE recommendations, Data Integration users can disable recommendations for individual mappings in the mapping designer.

If you create a mapping task that is based on an elastic mapping, you can enable CLAIRE recommendations to use CLAIRE Tuning.

Enterprise Data Catalog integration properties

If your organization uses Data Accelerator for Azure or data catalog discovery in Data Integration, you can configure Enterprise Data Catalog integration properties for your organization and sub-organizations. Configure Enterprise Data Catalog integration properties so that users can use catalog assets in mappings, synchronization tasks, file ingestion tasks, and Azure data sync tasks.

The Enterprise Data Catalog integration properties that you configure for the organization apply to the data catalog searches that all users in the organization perform. If your organization includes sub-organizations, you can configure different Enterprise Data Catalog integration properties for the parent organization and for each sub-organization.

The following table describes the Enterprise Data Catalog integration properties:

Property	Description
Catalog URL	URL of the Enterprise Data Catalog Service. Use the following format: `http://<fully qualified host name>:<port>` Do not append `/ldmcatalog` at the end of the URL.
Runtime environment	Name of the Secure Agent group that is used to read data from Enterprise Data Catalog. The agents in the group that you select must be able to communicate with Enterprise Data Catalog. Therefore, the Enterprise Data Catalog host must be in the same network as the agent machines or it must have the appropriate ports open for communication.
User name	Enterprise Data Catalog user account that the Secure Agent uses to access Enterprise Data Catalog. This user account must have privileges to view and search for objects in Enterprise Data Catalog and to perform functions using the Enterprise Data Catalog REST API.

Property	Description
Password	Password for the Enterprise Data Catalog user account.
Show the data catalog	Shows and hides the **Data Catalog** page in Data Integration.

Sub-organizations

If your organization has the Organization Hierarchy license, you can create one or more sub-organizations within your organization. Create sub-organizations to represent different business environments within your company. For example, you might create separate sub-organizations to represent your development, testing, and production environments.

When you create a sub-organization, the organization that you use to create a sub-organization becomes the parent organization. Each sub-organization can have only one parent, and it cannot contain another sub-organization.

The Organization Hierarchy license controls the number of sub-organizations that you can create. To increase this number, contact Informatica Global Customer Support.

Creating sub-organizations provides the following advantages:

You can manage sub-organization licenses individually or you can automatically synchronize them with the parent organization licenses.

Each sub-organization inherits all feature, connector, and custom licenses from the parent organization except for the Organization Hierarchy license and bundle licenses.

When you manage licenses individually, administrators for the parent organization can disable, enable, and modify the expiration dates for the licenses that the sub-organizations inherit. They configure the licenses separately for each sub-organization. Therefore, disabling a license in one sub-organization does not disable the license in other sub-organizations.

Alternatively, if you have the appropriate license, you can automatically synchronize sub-organization licenses with the parent organization. When this license is enabled, each time a license is changed in the parent organization, all sub-organizations inherit the license change.

You can manage users and assets separately.

Each sub-organization has its own set of users and assets.

Users whom you create in a sub-organization are unique to the sub-organization. They cannot log in to the parent organization or to other sub-organizations. Only administrators in the parent organization and users in the parent organization that have sub-organization access privileges can access the parent organization and all sub-organizations.

Assets such as mappings and tasks are also unique within an organization. Assets are not shared among sub-organizations or between the parent organization and any sub-organization. If you want to migrate an asset between organizations, export the asset from one organization and import it into a different organization.

You can share runtime environments.

Administrators in the parent organization can share Secure Agent groups with the sub-organizations. When you share Secure Agent groups, users in the sub-organizations can run tasks on the Secure Agents within the group.

You can switch between organizations without logging in to each one.

Users in the parent organization that have privileges to view sub-organizations can switch between organizations without logging out and logging back in to Informatica Intelligent Cloud Services.

Sub-organizations example

You need separate environments for synchronization task development, testing, and production to ensure that the tasks are tested before they are run in the production environment.

Log in to the parent organization as an administrator and create a sub-organization for task development, a sub-organization for testing, and a sub-organization for production. Add users to the sub-organizations, and ensure that the Data Synchronization license is enabled for each organization.

When task development is complete, export the tasks from the development sub-organization, and then log in to the testing sub-organization and import the tasks. When testing is complete, export the tasks from the testing sub-organization, and then log in to the production sub-organization and import the tasks.

Adding or removing a sub-organization

To add a sub-organization, you can either create a new sub-organization or link existing organizations. To remove a sub-organization, you can unlink organizations or delete the sub-organization.

You can add a sub-organization in either of the following ways:

- Create a sub-organization. Log in to the organization that you want to be the parent organization and create a sub-organization.
- Link existing organizations. The organization that you link from becomes the parent organization and the organization that you link to becomes a sub-organization.

You can remove a sub-organization in either of the following ways:

- Unlink an existing sub-organization from its parent organization.
- Delete the sub-organization. When you delete an organization, you delete all of the data associated with the organization.

Creating a sub-organization

The administrator of the parent organization can create a sub-organization.

1. Log in to the organization that you want to be the parent organization.
2. Open Administrator and select **Organization**.
3. Open the **Sub-Organizations** page and click **New Sub-Organization**.
4. Enter the properties for the sub-organization and click **Save**.

After you create a sub-organization, verify the licenses, and configure runtime environments, user accounts, and connections so that other people can use it.

Linking organizations

You can create a sub-organization by linking existing organizations. The organization that you link from becomes the parent organization and the organization that you link to becomes a sub-organization.

Before you link an organization, you need the organization ID for the organization that you want to link. You can find this information on the Organization page.

Note: If you link a sub-organization that has a license that the parent organization does not have, then the sub-organization loses the license.

You can link an organization if all of the following conditions apply:

- You have a user account with the organization.
- The organization is not the parent of another organization or a sub-organization of another organization.
- You are the administrator of the parent organization with the Organization Hierarchy license.
- The organization that you want to link as a sub-organization does not have the Organization Hierarchy license.

You can later unlink the organizations.

To link organizations:

1. Log in to the organization that you want to be the parent organization.
2. Open Administrator and select **Organization**.
3. Open the **Sub-Organizations** page and click **Link Sub-Organization**.
4. In the **Link Sub-Organization** dialog box, enter the following information:
   - The organization ID for the organization you want to link.
   - The user name and password of an administrator in the organization that you want to set up as a sub-organization.
5. Click **Link Sub-Organization** to link the organization.

   The organization displays on the **Sub-organizations** page.

Unlinking a sub-organization

You can unlink a sub-organization from your parent organization. After you unlink an organization, update the unlinked organization with the required licenses.

You can unlink a sub-organization if the following conditions apply:

- You have an administrator account with the sub-organization you want to unlink.
- You are the administrator of the parent organization with the Organization Hierarchy license.
- No asset in the sub-organization that you want to unlink uses a shared Secure Agent group as the runtime environment. If any asset in the sub-organization uses a shared Secure Agent group as the runtime environment, update the asset to use a different runtime environment before you unlink the sub-organization.

To unlink a sub-organization:

1. Log in to the parent organization.
2. Open Administrator and select **Organization**.
3. Open the **Sub-organizations** page.
4. Expand the Actions menu for the sub-organization that you want to unlink and select **Unlink**.

5. In the **Unlink** dialog box, enter the user name and password of a user in the sub-organization with the Admin role.
6. Click **Unlink**.

   The organizations are no longer linked.

Deleting a sub-organization

You can delete a sub-organization. When you delete a sub-organization, you delete all of the associated data.

You can delete a sub-organization if you are the administrator of the parent organization.

1. Log in to the parent organization.
2. Open Administrator and select **Organization**.
3. Open the **Sub-organizations** page.
4. Expand the Actions menu for the sub-organization that you want to unlink and select **Delete**.

Disabling or enabling a sub-organization

If you are the administrator for a parent organization, you can disable or enable a sub-organization.

When you create a sub-organization, the sub-organization is enabled by default. You might want to disable a sub-organization if you have a separate license agreement with the sub-organization and the license agreement expires. You can re-enable the sub-organization after you disable it.

You can disable or enable a sub-organization even if the sub-organization administrator blocks parent organization access to the sub-organization.

You can perform the following actions:

Disable a sub-organization

When you disable a sub-organization, the organization exists, but sub-organization users cannot log in to the sub-organization or access it through the REST API. Scheduled jobs in the sub-organization do not run.

Enable a sub-organization

When you enable a sub-organization, sub-organization users can log in to the sub-organization and access assets and perform tasks based on their user roles. Users with the appropriate privileges can access the sub-organization through the REST API. Scheduled jobs resume according to their schedules.

Disable or enable a sub-organization on the **Sub-Organizations** tab of the **Organizations** page. In the **Actions** menu for the sub-organization, select **Disable** or **Enable**.

Switching to a different organization

If you are an administrator in a parent organization or a user in a parent organization that has privileges to view sub-organizations, you can switch among organizations. You do not have to log out and log back in to Informatica Intelligent Cloud Services.

To switch to a different organization:

- From the **Organization** menu in the upper right corner, select the organization that you want to view.

Denying parent organization access to a sub-organization

If you are the administrator for a sub-organization, you can deny parent organization access to the sub-organization.

When you deny access to the sub-organization, users in the parent organization cannot switch from the parent organization to the sub-organization. Users in the parent organization with the appropriate privileges can make only the following changes to the sub-organization:

- Enable and disable the sub-organization
- Update the sub-organization licenses
- Edit the sub-organization properties such as the organization description and CLAIRE recommendation preferences

To deny parent organization access to the sub-organization, log in to the sub-organization as an administrator. On the **Organization** page, enable the **Deny parent organization access to this sub-organization** option.

Add-on connectors in sub-organizations

To use an add-on connector in a sub-organizations, you must install the connector in the parent organization. You cannot install add-on connectors in a sub-organization.

Sub-organizations inherit all connector licenses from the parent organization. If a sub-organization should not use a specific connector, disable the connector license for the sub-organization.

Exporting and importing assets in sub-organizations

Export and import assets in a sub-organization in the following ways:

- Log in to the sub-organization and export or import assets from within the sub-organization.
- Parent organization administrators can log in to the parent organization, switch to the sub-organization, and import or export Data Integration assets.

CHAPTER 3

Licenses

Licenses determine the Informatica Intelligent Cloud Services subscription level for the organization and provide access to Informatica Intelligent Cloud Services features, connectors, and bundles.

As an administrator, you can review the licenses that are set up for your organization, verify license expiration dates, and check job limits and usage. You can also manage sub-organization licenses and view job limits and usage for your sub-organizations.

License categories

Licenses are categorized as edition licenses, connector licenses, and custom licenses.

The following license categories are available:

Edition licenses

Edition licenses control the Informatica Intelligent Cloud Services features that you can use. Feature licenses provide access to data integration tasks such as mapping tasks, replication tasks, and synchronization tasks. They also provide access to components such as business services and saved queries and to features such as fine-grained security and Salesforce connectivity.

Connector licenses

Connector licenses provide connectivity to entities such as Amazon Redshift, Microsoft SQL Server, and Oracle.

Custom licenses

Custom licenses are licenses that are not part of an edition. They provide access to features, packages, or bundles. If your organization uses a custom license that provides access to a feature that is also included in an edition license, the terms of the custom license override the terms of the edition license.

License types

When you create an organization, Informatica Intelligent Cloud Services assigns the organization a license type for each licensed edition.

Informatica Intelligent Cloud Services uses the following types of licenses:

Trial

You can use the edition free of charge for a 30-day period. At the end of the trial period, you can subscribe to the edition. A trial subscription might provide limited access to the features, connectors, and packages that are associated with the license.

Subscription

You can use the licensed edition for the duration of the contract period. Near the end of the contract period, Informatica Intelligent Cloud Services indicates that the contract is about to expire. Renew the contract to continue to use the edition.

Free subscription

You can use the synchronization task free of charge. A free subscription might provide limited access to the features of the synchronization task.

Sub-organization licenses

A sub-organization has licenses that are maintained by the parent organization. If a sub-organization requires a license that does not belong to the parent organization, contact Informatica Global Customer Support to obtain the license for the parent organization.

When you create a sub-organization, each sub-organization inherits licenses from the parent organization as custom licenses. The sub-organization inherits all licenses except for the following licenses:

- The Organization Hierarchy license
- Bundle licenses. To use a bundle in the sub-organization, a user in the sub-organization must install the bundle.

You can manage sub-organization licenses in the following ways:

Manage the licenses individually

When you manage licenses individually, administrators for the parent organization can disable, enable, and shorten the expiration dates for the inherited licenses. They manage the licenses separately for each sub-organization. Sub-organization administrators can view licenses but cannot change them.

This is the default option.

Automatically synchronize sub-organization licenses with the parent organization

If you have the appropriate license, you can automatically synchronize sub-organization licenses with the parent organization. When this license is enabled, each time a license is changed in the parent organization, all sub-organizations inherit the license change.

You might want to enable license synchronization when your organization has many sub-organizations and the sub-organizations have the same licenses.

If license synchronization is not enabled for your organization, then you must manage sub-organization licenses individually.

Note: If you link a sub-organization that has a license that the parent organization does not have, the sub-organization loses the license.

Editing sub-organization licenses

You can edit sub-organization licenses if you are an administrator in the parent organization and if license synchronization between the parent organization and sub-organizations is not enabled. You can edit sub-organization licenses from within the parent organization or from within the sub-organization.

1. Log in to the parent organization.
2. To edit licenses from within the parent organization:
   a. Open Administrator and select **Organization**.
   b. Click **Sub-Organizations**.
   c. Select the sub-organization for which you want to edit licenses.
   d. Click **Licenses**.
3. To edit licenses from within the sub-organization:
   a. From the **Organization** menu in the upper right corner, select the sub-organization for which you want to edit licenses.
   b. Open Administrator and select **Licenses**.
4. Select licenses to enable features, and clear licenses to disable features.
5. Optionally, modify expiration dates.

   You can shorten, but not extend, license expiration dates for a sub-organization.
6. Click **Save**.

Synchronizing licenses with the parent organization

You can automatically synchronize sub-organization licenses with the parent organization. Each time a license is changed in the parent organization, all sub-organizations inherit the license change.

To enable license synchronization, contact Informatica Global Customer Support and request the license for this feature. When the license is enabled for the parent organization, license synchronization with the sub-organizations happens automatically. The parent organization administrator does not have to take any action to synchronize the licenses.

Note: When the license for this feature is enabled, you cannot edit sub-organization licenses individually.

When the license for this feature is enabled and you disable a sub-organization, the sub-organization loses its license settings. When you re-enable the sub-organization, the sub-organization inherits all license settings from the parent organization.

License synchronization between a parent organization and sub-organizations does not affect the license meter counts in the sub-organizations.

License expiration

When a license expires, you cannot access the features, connectors, or packages that are associated with the license. Scheduled jobs that are associated with the license are also disabled. If all licenses for the organization expire, you cannot log in to Informatica Intelligent Cloud Services.

You can review the expiration date for licenses on the **Licenses** page in Administrator. To extend a license, contact Informatica Global Customer Support. After you extend a license, you can access the associated features, connectors, and packages, and the scheduled jobs resume processing.

CHAPTER 4

Ecosystem single sign-on

Informatica Intelligent Cloud Services enables single sign-on capability for Microsoft Azure users. This allows Microsoft Azure users to sign in to Informatica Intelligent Cloud Services without having to enter their login information again.

When you create your organization through Microsoft Azure, you can configure some single sign-on properties for Microsoft Azure users on the **Ecosystem SSO** page.

Note: The ecosystem single sign-on properties that you configure for Microsoft Azure are different than the SAML single sign-on properties that you configure to enable single-sign on from a third-party identity provider. To configure SAML single sign-on for your organization, see Chapter 5, “SAML single sign-on” on page 30.

The following image shows the **Ecosystem SSO** page:

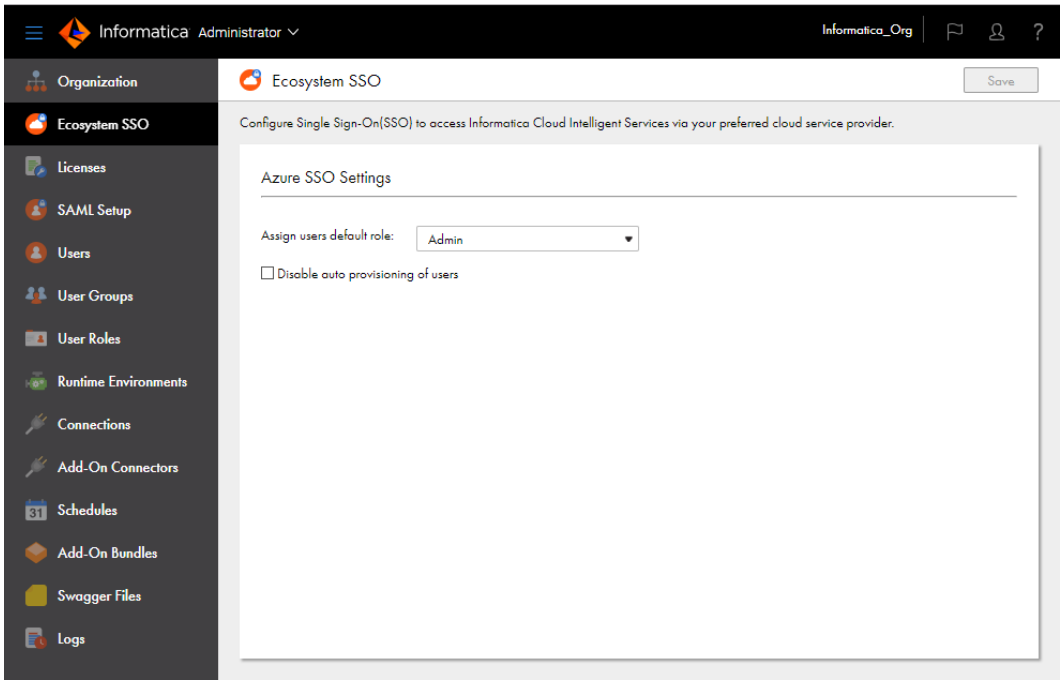

You can configure the following properties for Microsoft Azure users:

Assign users default role

When a Microsoft Azure user signs in to your organization for the first time, Informatica Intelligent Cloud Services adds the user to your organization and assigns the user a default role. By default, Informatica Intelligent Cloud Services assigns the user the Admin role.

You can change the default role to a different role such as the Designer role. To change the default user role, select a different role in the **Assign users default role** list.

Note: If you want Microsoft Azure users to be able to download, install, and register a Secure Agent, assign them the Admin or Designer role. You can also assign users a custom role that has privileges to create, read, and update Secure Agents.

Disable auto-provisioning of users

By default, the first time a Microsoft Azure user signs in to Data Accelerator for Azure, Informatica Intelligent Cloud Services adds the user to your organization. This process is called auto-provisioning.

You can enable or disable the auto-provisioning of Microsoft Azure users. To do this, enable or disable the **Disable auto provisioning of users** option.

Note: If you disable auto-provisioning, you must create each user on the **Users** page. If you want the user to be able to use single sign-on from Microsoft Azure, you must also set the **Authentication** field on the user details page to **Azure SSO**.

CHAPTER 5

SAML single sign-on

You can enable single sign-on (SSO) capability so that users can access their organization without the need to enter login information.

Single sign-on to Informatica Intelligent Cloud Services is based on the Security Assertion Markup Language (SAML) 2.0 web browser single sign-on profile. The SAML web browser single sign-on profile consists of the following entities:

Identity provider

An entity that manages authentication information and provides authentication services through the use of security tokens.

Service provider

An entity that provides web services to principals, for example, an entity that hosts web applications. Informatica Intelligent Cloud Services is a service provider.

Principal

An end user who interacts through an HTTP user agent.

SAML 2.0 is an XML-based protocol that uses security tokens that contain assertions to pass information about a principal between an identity provider and a service provider. An assertion is a package of information that supplies statements made by a SAML authority.

When a user enters the Informatica Intelligent Cloud Services single sign-on URL in a browser, the following process occurs:

1. Informatica Intelligent Cloud Services sends a SAML authentication request to the organization's identity provider.
2. The identity provider confirms the user's identity and sends a SAML authentication response to Informatica Intelligent Cloud Services.
3. When Informatica Intelligent Cloud Services receives the SAML authentication response from the identity provider, Informatica Intelligent Cloud Services establishes the user session and logs the user into Informatica Intelligent Cloud Services.
4. When a user logs out of Informatica Intelligent Cloud Services or the session times out, Informatica Intelligent Cloud Services sends a SAML logout request to the identity provider.
5. The identity provider terminates the user session on the identity provider side.

You can find more information about SAML on the Oasis web site: https://www.oasis-open.org

SAML single sign-on requirements

To set up SAML single sign-on for an Informatica Intelligent Cloud Services organization, the system must use an appropriate identity provider. You must also have the appropriate license.

To set up SAML single sign-on for an organization, ensure that the following requirements are met:

- The system must use a SAML 2.0-based identity provider.

  Common identity providers include Microsoft Active Directory Federation Services (AD FS), Okta, SSOCircle, OpenLDAP, and Shibboleth. The identity provider must be configured to use either the DSA-SHA1 or RSA-SHA1 algorithm to generate the signature.
- The Informatica Intelligent Cloud Services organization must have the SAML based Single Sign-On license.
- You must have access to the organization as an organization administrator to set up single sign-on.

Single sign-on restrictions

There are some restrictions for SAML single sign-on access to Informatica Intelligent Cloud Services.

The following restrictions apply to SAML single sign-on access:

- If your license with the identity provider expires, you cannot access Informatica Intelligent Cloud Services through single sign-on.
- If the identity provider is down or Informatica Intelligent Cloud Services servers cannot reach it, users cannot log in to Informatica Intelligent Cloud Services through single sign-on.
- If the identity provider certificate used for SAML single sign-on to Informatica Intelligent Cloud Services expires, users cannot access Informatica Intelligent Cloud Services through single sign-on.
- If your organization uses trusted IP address ranges, users cannot log in to Informatica Intelligent Cloud Services from an IP address that is not within the trusted IP address ranges.

User management with SAML single sign-on

The following rules apply to users and user accounts when you enable SAML single-sign on for Informatica Intelligent Cloud Services:

- Informatica Intelligent Cloud Services stores user information that passes from the identity provider such as first name and email address in the Informatica Intelligent Cloud Services repository.
- You can create a regular user account with credentials in Informatica Intelligent Cloud Services after you enable an organization for single sign-on, and the user credentials are saved in the Informatica Intelligent Cloud Services repository. However, the user must log in to Informatica Intelligent Cloud Services directly instead of using single sign-on.
- If you delete a user from Informatica Intelligent Cloud Services, the user is deleted from the Informatica Intelligent Cloud Services repository. The user is not deleted from the identity provider.

SAML single sign-on configuration for Informatica Intelligent Cloud Services

Informatica Intelligent Cloud Services and your identity provider exchange configuration information when you set up single sign-on.

Informatica Intelligent Cloud Services requires identity provider metadata to send authentication requests to the identity provider. The identity provider requires service provider metadata from Informatica Intelligent Cloud Services to send authentication responses to Informatica Intelligent Cloud Services.

SAML and Informatica Intelligent Cloud Services attributes such as user roles need to be mapped so that Informatica Intelligent Cloud Services can consume the data passed in authentication responses. After you configure single sign-on settings in Informatica Intelligent Cloud Services, pass the Informatica Intelligent Cloud Services service provider metadata to your identity provider.

To configure single sign-on for Informatica Intelligent Cloud Services, complete the following tasks:

1. Configure the SAML identity provider and service provider settings, and map SAML attributes and user roles to Informatica Intelligent Cloud Services attributes and user roles in Informatica Intelligent Cloud Services.
2. Download the Informatica Intelligent Cloud Services service provider metadata from Informatica Intelligent Cloud Services, and deliver the metadata and the Informatica Intelligent Cloud Services single sign-on URL for your organization to your SAML identity provider administrator.

Configuring provider settings and mapping attributes

Configure SAML single sign-on settings and map SAML attributes on the **SAML Setup** page.

1. Log in to Informatica Intelligent Cloud Services as an organization administrator.
2. In Administrator, select **SAML Setup**.
3. On the **SAML Setup** page, configure the following properties:
   - Identity provider properties
   - Service provider properties
   - SAML attribute mapping properties
   - SAML role mapping properties
4. Click **Save**.

   Informatica Intelligent Cloud Services generates the service provider metadata file. Informatica Intelligent Cloud Services also generates a unique token for your organization and saves the token to the Informatica Intelligent Cloud Services repository. The single sign-on URL for your organization includes the token. For example:

   ```
   https://dm-us.informaticacloud.com/ma/sso/<organization token>
   ```

After you save your changes on the **SAML Setup** page, download the service provider metadata, and send it to your identity provider along with the Informatica Intelligent Cloud Services single sign-on URL.

Identity provider properties

Define SAML identity provider properties on the **SAML Setup** page.

If you have an identity provider XML file, you can upload the file to populate some of the properties. Informatica Intelligent Cloud Services can parse and extract most of the data from the XML file. However, you might need to enter certain fields manually such as the name identifier format.

The following table describes identity provider configuration properties:

Property	Description
Use Identity Provider File	The identity provider XML file populates many of the properties on the **SAML Setup** page. To use an identity provider XML file to define identity provider properties, click **Browse**, and navigate to the identity provider XML file.
Disable auto provisioning of users	Disables auto provisioning of SAML users. When a new SAML user logs in to Informatica Intelligent Cloud Services for the first time, the user will not be added to the organization in Informatica Intelligent Cloud Services.
Issuer	The entity ID of the identity provider, which is the unique identifier of the identity provider. The Issuer value in all messages from the identity provider to Informatica Intelligent Cloud Services must match this value. For example: `<saml:Issuer>http://idp.example.com</saml:Issuer>`
Single Sign-On Service URL	The identity provider's HTTP-POST SAML binding URL for the SingleSignOnService, which is the SingleSignOnService element's location attribute. Informatica Intelligent Cloud Services sends login requests to this URL.
Single Logout Service URL	The identity provider's HTTP-POST SAML binding URL for the SingleLogoutService, which is the SingleLogoutService element's location attribute. Informatica Intelligent Cloud Services sends logout requests to this URL.
Signing Certificate	Base64-encoded PEM format identity provider certificate that Informatica Intelligent Cloud Services uses to validate signed SAML messages from the identity provider. **Note:** The identity provider signing algorithm must be either DSA-SHA1 or RSA-SHA1.
Use signing certificate for encryption	Uses the public key in your signing certificate to encrypt logout requests sent to your identity provider when a user logs out from Informatica Intelligent Cloud Services.
Encryption Certificate	Base64-encoded PEM format identity provider certificate that Informatica Intelligent Cloud Services uses to encrypt SAML messages sent to the identity provider. Applicable if you do not enable use of the signing certificate for encryption.
Name Identifier Format	The format of the name identifier in the authentication request that the identity provider returns to Informatica Intelligent Cloud Services. Informatica Intelligent Cloud Services uses the name identifier value as the Informatica Intelligent Cloud Services user name. The name identifier cannot be a transient value that can be different for each login. For a particular user, each single sign-on login to Informatica Intelligent Cloud Services must contain the same name identifier value. To specify that the name identifier is an email address, the Name Identifier Format is as follows: `urn:oasis:names:tc:SAML:1.1:nameidformat:emailAddress`

Property	Description
Logout Service URL (SOAP Binding)	The identity provider's SAML SOAP binding URL for the single logout service. Informatica Intelligent Cloud Services sends logout requests to this URL.
Logout Page URL	The landing page to which a user is redirected after the user logs out of Informatica Intelligent Cloud Services. Informatica Intelligent Cloud Services redirects the logged out user to the landing page in the following ways: - If you specify a logout page URL, Informatica Intelligent Cloud Services redirects the user to this URL after logout. - If you do not specify a logout page URL, Informatica Intelligent Cloud Services redirects the user to a default logout page.

Service provider properties

Define the Informatica Intelligent Cloud Services service provider properties on the **SAML Setup** page.

The following table describes service provider properties:

Property	Description
Informatica Cloud Platform SSO	Displays the single sign-on URL for your organization. This URL is automatically generated by Informatica Intelligent Cloud Services.
Clock Skew	Specifies the maximum permitted time between the time stamps in the SAML response from the identity provider and the Informatica Intelligent Cloud Services clock.
Name Identifier value represents user's email address	If selected, Informatica Intelligent Cloud Services uses the name identifier as the email address.
Sign authentication requests	If selected, Informatica Intelligent Cloud Services signs authentication requests to the identity provider.
Sign logout requests sent using SOAP binding	If selected, Informatica Intelligent Cloud Services signs logout requests sent to the identity provider.
Encrypt name identifier in logout requests	If selected, Informatica Intelligent Cloud Services encrypts the name identifier in logout requests. **Note:** Verify that the identity provider supports decryption of name identifiers.

SAML attribute mapping properties

User login attributes such as name, email address, and user role are included in the authentication response from the identity provider to Informatica Intelligent Cloud Services. Map the Informatica Intelligent Cloud Services user fields to corresponding SAML attributes on the **SAML Setup** page.

The following table describes SAML attribute mapping properties:

Property	Description
Use friendly SAML attribute names	If selected, uses the human-readable form of the SAML attribute name which might be useful in cases in which the attribute name is complex or opaque, such as an OID or a UUID.
First Name	SAML attribute used to pass the user first name.
Last Name	SAML attribute used to pass the user last name.
Job Title	SAML attribute used to pass the user job title.
Email Addresses	SAML attribute used to pass the user email addresses.
Emails Delimiter	Delimiter to separate the email addresses if multiple email addresses are passed.
Phone Number	SAML attribute used to pass the user phone number.
Time Zone	SAML attribute used to pass the user time zone.
User Roles	SAML attribute used to pass the user assigned user roles.
Roles Delimiter	Delimiter to separate the roles if multiple roles are passed.

SAML role mapping properties

Map SAML role names to Informatica Intelligent Cloud Services role names. You can map multiple SAML role names to a single Informatica Intelligent Cloud Services role. Define the SAML role mapping properties on the **SAML Setup** page.

The following table describes SAML role mapping properties:

Property	Description
Informatica Intelligent Cloud Services role	The SAML role equivalent for the Informatica Intelligent Cloud Services role. If you need to enter more than one role, use a comma to separate the roles.
Default Role	Default role to use if the SAML authentication response does not include the SAML user roles attribute.
Default User Group	Default user group for single sign-on users.

Downloading the service provider metadata

The identity provider requires the SAML service provider metadata and Informatica Intelligent Cloud Services URL to complete the SAML single sign-on setup process. After Informatica Intelligent Cloud Services generates the service provider metadata file, deliver the file and the Informatica Intelligent Cloud Services URL to the identity provider.

1. On the **SAML Setup** page, click **Download Service Provider Metadata**.

   The service provider metadata file is downloaded to your machine.
2. In the **Information** dialog box, note the URL for single sign-on access to your Informatica Intelligent Cloud Services organization.
3. Click **OK** to close the **Information** dialog box.
4. Send the metadata file and the Informatica Intelligent Cloud Services single sign-on URL to your identity provider administrator.

CHAPTER 6

Metering

You can view metering information for your organization and sub-organizations. View metering information on the **Metering** page.

Meters display the amounts of computing resources that your organization uses such as the total data volume for Mass Ingestion Streaming. Meters also display the job limits set through your organization's licenses.

The **Metering** page displays information in the following views based on your organization's licenses:

Dashboard view

If you have the appropriate licenses, the **Metering** page displays information in a dashboard view.

In the dashboard view, the summary area displays a summary of the remaining computing resources for the month. You can also display a table of all meters that apply to your organization, which are determined by your organization's editions.

The detail area displays monthly usage statistics for Mass Ingestion Streaming data volume usage. You can display a detailed chart of the data volume usage.

License Metrics view

If you do not have the licenses required to see the dashboard view, the **Metering** page displays license metrics in a table. The table lists all meters that apply to your organization, which are determined by your organization's editions.

Viewing license metrics

You can view a table of all meters used in your organization and download a report that shows the metering usage. View the table and download the report from the License Metrics view of the **Metering** page.

To open the License Metrics view from the dashboard view, click **All Meters** in the Metrics Summary This Month area. If you do not have the licenses required to see the dashboard view, the License Metrics view is displayed when you open the **Metering** page.

The meters that appear in the License Metrics view are determined by the editions that your organization has. Your organization might also be assigned custom meters. Metering information might not be available for every edition.

If your organization has multiple editions or uses custom meters, a meter might be displayed multiple times with different limits. In this case, the least restrictive limit applies. For example, if one edition has a limit of 500 synchronization jobs per day and another edition has a limit of 100 synchronization jobs per day, the 500 job per day limit applies. The **In Effect** column indicates which limit applies.

The License Metrics view displays the following information for each meter:

Property	Description
Edition	Name of the edition that is associated with the meter.
Service	Service to which the meter applies.
Metering	Name of the meter. For example, the number of synchronization jobs per day, the number of rows processed by mapping jobs per month, or the total number of replication jobs.
Limit	Numeric limit such as the maximum number of jobs or processed rows. The limit applies to the parent organization and to each sub-organization. For example, if the limit is 100 jobs per day, users in the parent organization can run 100 jobs per day, and users in each sub-organization can also run 100 jobs per day. If this field displays -1, there is no limit.
Used	The actual number units consumed, such as the number of jobs run or compute hours used, in the organization or sub-organization during the metering period.
Percent Used	The percentage of units consumed in the organization or sub-organization during the metering period.
In Effect	Indicates whether the meter is in effect for the organization or sub-organization.

Meter definitions

The meters that appear in the License Metrics view of the **Metering** page are determined by the editions that your organization has.

The following table describes the meters that might be in effect based on your editions:

Meter	Definition
Total number of agents	Total number of Secure Agents, including agents that are stopped. Does not include the Informatica Cloud Hosted Agent.
Total number of connections	Total number of connections.
Total number of projects	Total number of projects.
Total number of folders	Total number of folders.
Daily/monthly incoming API request maximum	Number of API access requests per day or per month.
Number of mapping jobs per day/month	Number of mapping jobs per day or per month.*
Total number of mapping jobs	Total number of mapping jobs.*
Number of rows processed by mapping jobs per day/month	Number of rows processed by mapping jobs per day or per month.*
Total number of rows processed by mapping jobs	Total number of rows processed by mapping jobs.*
Number of masking jobs per day/month	Number of masking jobs per day or per month.

Meter	Definition
Total number of masking jobs	Total number of masking jobs.
Number of rows processed by masking jobs per day/month	Number of rows processed by masking jobs per day or per month.
Total number of rows processed by masking jobs	Total number of rows processed by masking jobs.
Number of PowerCenter jobs per day/month	Number of PowerCenter jobs per day or per month.
Total number of PowerCenter jobs	Total number of PowerCenter jobs.
Number of rows processed by PowerCenter jobs per day/month	Number of rows processed by PowerCenter jobs per day or per month.
Total number of rows processed by PowerCenter jobs	Total number of rows processed by PowerCenter jobs.
Number of replication jobs per day/month	Number of replication jobs per day or per month.
Total number of replication jobs	Total number of replication jobs.
Number of rows processed by replication jobs per day/month	Number of rows processed by replication jobs per day or per month.
Total number of rows processed by replication jobs	Total number of rows processed by replication jobs.
Number of data (in GB) processed by the streaming ingestion task per month	Number of gigabytes (GB) processed by streaming ingestion jobs per month.
Number of synchronization jobs per day/month	Number of synchronization jobs per day or per month.
Total number of synchronization jobs	Total number of synchronization jobs.
Number of rows processed by synchronization jobs per day/month	Number of rows processed by synchronization jobs per day or per month.
Total number of rows processed by synchronization jobs	Total number of rows processed by synchronization jobs.
Number of sub-organizations	Number of sub-organizations.
Number of state sync jobs per day	Number of state synchronization jobs per day. State synchronization jobs include the fetchState and loadState jobs that you run through the REST API.
Number of user-management create requests	Number of user management creation requests. User management creation requests include requests to create users, user groups, and custom roles.

Meter	Definition
Total compute hours used by elastic cluster nodes	Total compute hours used by elastic cluster nodes. **Note:** If your organization uses Data Integration Elastic, metering is calculated based on the processing power that each node provides in an elastic cluster or compute hours. Metering begins when a node is added to the cluster. Metering ends when the node is removed from the cluster or the cluster is stopped.
Total Serverless units used	Total number of serverless compute units used to run tasks.
* If your organization uses Data Accelerator for Azure, this meter includes Azure data sync jobs because each Azure data sync job runs an underlying mapping.	

Metering serverless compute units

When you view the total serverless compute units used, the meter is based on the number of serverless compute units that your organization uses to run tasks.

When the serverless runtime environment runs a task, the environment creates a virtual machine with resources based on the number of compute units that the task requests.

The minimum task duration is five minutes. If the task completes in less than five minutes, the serverless runtime environment consumes five minutes of compute units. After five minutes, compute units are consumed by the second.

If you cancel the job, the number of consumed compute units is the greater of the following values:

- The time that the job was running before it was canceled
- Five minutes

Guidelines for Data Integration Elastic

In Data Integration Elastic, the serverless runtime environment creates an elastic cluster that contains one worker node with resources based on the number of serverless compute units that the task requests. If you run another task, the environment reuses idle worker nodes or adds worker nodes to the cluster to reserve additional resources.

Metering begins when the task starts running and ends when the task is complete. Metering doesn't include the time to compile the job or the time to start the cluster.

Metering doesn't take effect if the job fails before the cluster has been created, such as when the job fails to compile, the cluster fails to start, or you cancel the job before the cluster starts.

Metering usage reports

If your organization uses Data Integration Elastic, the Mass Ingestion service, or a serverless runtime environment, you can download a metering usage report.

A metering usage report contains the following information:

- For Data Integration Elastic, the report contains details about the compute hours for nodes in elastic clusters. You can export the details to a CSV file.
- For Mass Ingestion, the report contains details about the volume of data ingested by Mass Ingestion Streaming jobs.
- For a serverless runtime environment, the report contains details about the number of serverless compute units that were requested and consumed to run tasks.

If your organization does not use Data Integration Elastic, the Mass Ingestion Streaming, or a serverless runtime environment, no metering usage reports are available.

Downloading a metering usage report

Download a metering usage report from the License Metrics view of the **Metering** page.

1. In Administrator, select **Metering**.
2. If you see the dashboard view, click **All Meters** in the Metrics Summary This Month area to open the License Metrics view.
3. Click **Export to File** and select **Metering Usage Details**.
4. Select the service or product feature and the date range that you want to view, and then click **Export**.

Viewing usage details

You can view detailed statistics for Mass Ingestion Streaming data volume usage. To view the detailed statistics, in the dashboard view of the **Metering** page, click **Detail Chart** or click the graph in the Data Volume Usage This Month area.

The following image shows an example of the details page for data volume usage over the last 90 days:

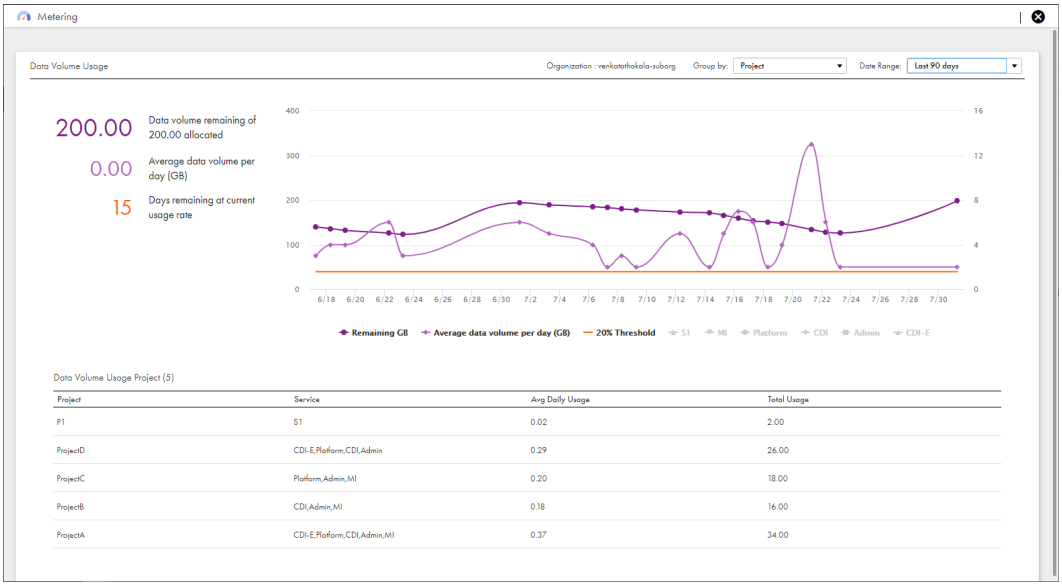

You can customize the page in the following ways:

- If your organization has sub-organizations, you can view usage details for the parent organization or for a sub-organization.
- You can change the grouping. For example, you can view usage details by project, runtime environment, or both.
- You can change the date range.
- If the usage details apply to multiple services, you can update the graph to display the resource usage for each service. To update the graph to display the resource usage for a service, select the service name below the graph.

Note: In Mass Ingestion Streaming, if you change the runtime environment name or project of a streaming ingestion job that is in running state, metering information continues to appear with the old run-time environment name or in the old project of the job. Redeploy the streaming ingestion job to view the respective metering information in the new run-time environment name or project.

CHAPTER 7

Source control and service upgrade settings

You can configure source control settings and upgrade settings for Secure Agent services on the **Settings** page.

You can configure the following settings:

Source control settings

If your organization has the appropriate license, you can configure source control for your organization. Configure source control to enable version management for projects, folders, and assets. To configure source control, enable the **Enable Source Control** option, configure the type of repository access, and then configure the connection to the source control repository.

You can configure read/write or read-only access to the source control repository. To configure read/write access, enable the **Allow Push to Git Repository** option. To configure read-only access, disable the **Allow Push to Git Repository** option.

Upgrade settings for Secure Agent services

If a service that supports rolling upgrades encounters an error during an upgrade, you can specify whether to continue or stop upgrading the service.

You can also configure a restart schedule for services that need to be restarted after minor upgrades. To configure a restart schedule, select the day of the week and the time to perform the upgrades.

Source control configuration

You can configure source control for your organization to enable version management for projects, folders, and assets. When you configure source control, you can store versions of objects in a Git repository. Configure source control on the **Settings** page.

When you configure source control for an organization, users can apply source control to objects. Objects are not checked in automatically. Users can apply source control to individual assets or to all assets in a project or folder. For more information about applying source control to projects, folders, and assets, see the help system for the appropriate Informatica Intelligent Cloud Services service.

To configure source control for an organization, you must have the appropriate license.

You can configure source control for your organization in the following ways:

Configure read/write access to the source control repository.

When you configure read/write access, users in your organization can check in and check out objects, pull versions of objects, and revert objects to a previous version. Users must check out source-controlled objects to change them. Users check out objects exclusively, so one user cannot change an object that is checked out by another user. Users can change objects that are not source-controlled without checking them out.

You might want to configure read/write access for an organization in which you develop projects and assets.

Configure read-only access to the source control repository.

When you configure read-only access, users in your organization can pull versions of source-controlled objects from the repository. However, users cannot check out or check in objects. Users can make changes to projects, folders, and assets in the organization without checking them out.

You might want to configure read-only access for a test or production organization so that users can test or run the latest versions of assets.

Warning: When you configure read-only access, users can overwrite source-controlled objects. For example, user John pulls the most recent version of a source-controlled mapping and changes it. If another user pulls any version of the mapping later, John's changes are lost. Configure object permissions and user privileges carefully to prevent users from accidentally overwriting source-controlled assets in your organization.

You can change the repository access type. However, to change from read/write to read-only, you must first ensure that no objects are checked out. Informatica Intelligent Cloud Services does not allow you to change the repository access type from read/write to read-only if any objects are checked out.

You can also change the repository URL. To do this, you must first unlink all source-controlled assets. Informatica Intelligent Cloud Services does not allow you to change the repository URL if any assets are source-controlled.

If you want to disable source control after you configure it, unlink all objects from source control and then disable source control for the organization.

Source control configuration for sub-organizations

Configure source control for a sub-organization on the **Settings** page in the sub-organization. As a best practice, each sub-organization should use its own source control repository. Additionally, the source control repository for a sub-organization should be different than the source control repository for the parent organization.

Maintain different source control repositories so that users in one organization do not accidentally overwrite or change assets in another organization.

If you want the parent organization administrator to be able to perform source control operations in the sub-organization, configure the Git user account for the parent organization administrator to have access to the sub-organization's source control repository.

Repository access using OAuth

You can configure an organization to use OAuth authentication instead of personal access tokens to provide access to the source control repository. Configure OAuth authentication on the **Settings** page.

If you use a GitHub repository, you must have a GitHub access application installed on your repository that allows Informatica Intelligent Cloud Services to perform source control operations on the organization's

GitHub repository. If you don't have this application installed on your repository, you can install it from the **Settings** page.

Enabling source control for an organization

To enable source control for an organization, configure the type of access and connection to the source control repository on the **Settings** page. The source control repository that you select must contain a branch named "master."

1. On the **Settings** page in Administrator, click **Edit** in the Source Control area.
2. Enable the **Enable Source Control** option.
3. Configure the type of access to the source control repository:
   - To configure read/write access, enable the **Allow Push to Git Repository** option.
   - To configure read-only access, disable the **Allow Push to Git Repository** option.
4. Enter the repository URL, for example:

   ```
   https://github.com/MyGitUser/MyRepositoryName.git
   ```

   The repository URL must use the HTTPS protocol.
5. Optionally, to use OAuth to access the repository, enable **Allow OAuth access to Git**.

   If you use a GitHub repository, a Git access application that authorizes Informatica Intelligent Cloud Services access must be installed on your repository. To install the application, click **Install Git Access App**.
6. Click **Save**.

   Informatica Intelligent Cloud Services prompts you for your source control credentials to verify the repository connection. Informatica Intelligent Cloud Services does not store this information.

   If the connection is valid and you configure read/write access to the repository, Informatica Intelligent Cloud Services writes a small readme file to the repository to verify that it can push objects to the repository.

After you enable source control, all users that use source control must enter their source control credentials in their user settings. Users cannot see source control columns on the **Explore** page or perform source control actions until they enter their source control credentials. To enter source control credentials, click the **User** icon in the top right corner of the Informatica Intelligent Cloud Services window, and select **Settings**.

Changing the source control repository URL

To change the source control repository URL, you must first unlink all objects in the organization and then enter the new repository URL on the **Settings** page in Administrator. After you change the URL, all users that use source control must update their source control credentials in the user settings.

1. In each Informatica Intelligent Cloud Services service that uses the repository, unlink all objects from source control.
2. In Administrator, open the **Settings** page and click **Edit** in the Source Control area.
3. Verify that the **Enable Source Control** option is enabled.
4. Configure the type of access to the source control repository:
   - To configure read/write access, enable the **Allow Push to Git Repository** option.
   - To configure read-only access, disable the **Allow Push to Git Repository** option.
5. Enter the new repository URL, for example:

```
https://github.com/MyGitUser/MyRepositoryName.git
```

Tip: In GitHub, you can find the repository URL by opening the repository and selecting **Clone or download** > **Clone with HTTPS**.

The repository URL must use the HTTPS protocol.

6. Click **Save**.

   Informatica Intelligent Cloud Services prompts you for your source control credentials to verify the repository connection. Informatica Intelligent Cloud Services does not store this information.

   If the connection is valid and you configure read/write access to the repository, Informatica Intelligent Cloud Services writes a small readme file to the repository to verify that it can push objects to the repository.

After you change the source control repository URL, all users that use source control must update their source control credentials in the user settings. To update source control credentials, click the **User** icon in the top right corner of the Informatica Intelligent Cloud Services window, and select **Settings**.

Disabling source control for an organization

You can disable source control for an organization. Disabling source control breaks the link between your organization and the source control repository. It does not delete objects in the source control repository.

Before you can disable source control for a read-write organization, all assets must be unlinked.

1. In each Informatica Intelligent Cloud Services service that uses the repository, unlink all objects from source control.
2. In Administrator, disable source control:
   a. In Administrator, open the **Settings** page.
   b. Click **Edit** in the Source Control area.
   c. Disable the **Enable Source Control** option.
   d. Click **Save**.
3. Optionally, have users in the organization delete their source control credentials in their user settings:
   a. In the top right corner of the Informatica Intelligent Cloud Services window, click the **User** icon and select **Settings**.
   b. Clear the source control credentials.
   c. Click **Save**.

Configuring repository access

To work with source controlled objects, specify your GitHub or Azure DevOps Git repository credentials in Informatica Intelligent Cloud Services.

Your credentials can include a user name and a personal access token.

If your administrator has configured the organization's repository for OAuth access, you can enable OAuth access instead of providing a personal access token.

Personal access tokens must be configured to enable full control of private repositories. For information about generating a personal access token, see the GitHub or Azure DevOps Git help.

In Informatica Intelligent Cloud Services, perform the following steps to configure access to the repository:

1. Click the **User** icon in the top right corner of the Informatica Intelligent Cloud Services window and then select **Settings**.

2. Perform one of the following tasks:
   - Enter your repository credentials. For GitHub, enter your user name and personal access token. For Azure DevOps Git, enter your personal access token.
   - Enable OAuth access to the repository. If you have not already authorized access, a Git access app appears. Select to authorize access for Informatica Intelligent Cloud Services.
3. Click **Save**.

Source control best practices

To ensure that your organization configures and uses source control effectively, use the following guidelines as best practices.

Setup guidelines

Adhere to the following guidelines when you set up source control for your organization:

- Use different organizations for development, testing, staging, and production.

  When you maintain different organizations, you maintain isolation across environments so that changes in one environment do not affect other environments. For example, changes to assets in the testing environment are not accidentally deployed in the production environment.
- Configure development organizations with read/write access to the source control repository, and configure non-development organizations with read-only access to the source control repository.

  This ensures that only users in a development organization can make changes to assets. It also prevents users in a non-development environment from accidentally pushing changes to the source control repository.
- Ensure that only one development organization uses a particular source control repository.

  Maintaining separate repositories ensures that users in one organization do not accidentally change or overwrite assets in a different organization.
- When you enable source control for the organization, select an empty repository.

  Ensure that the repository does not contain a folder named "Explore" because Informatica Intelligent Cloud Services stores assets under the Explore folder in the Git repository.
- Do not share source control credentials among multiple Informatica Intelligent Cloud Services users.

  Separate credentials maintain security and make it easier to track which user made a particular change. Additionally, each user gets their own rate limit in GitHub.

Development guidelines

Adhere to the following guidelines as you develop and work with assets:

Guidelines for managing dependencies

Use the following guidelines to manage assets with dependencies:

- Create connections and runtime environments before you pull assets from the repository.

  When required connections and runtime environments exist in the target organization, you can run tasks immediately after you pull them from the repository.
- Ensure that reusable assets such as mappings and components are present in the repository before you use them.

Informatica Intelligent Cloud Services does not allow you to save an asset such as a mapping task when the dependent mapping does not exist in the organization.

- Avoid moving or renaming source-controlled assets that are used by other assets.

  If you move or rename a source-controlled asset, references to the asset can break.

Guidelines for checking in and checking out assets

Use the following guidelines when you check in and check out assets:

- Identify all dependencies before you check out reusable assets such as mappings, mapplets, and user-defined functions.

  Source control operations such as check out, check in, and pull do not automatically include dependent assets.

- When you need to update a reusable asset such as a mapping or component, check out the asset and all dependent assets.

  For example, when you need to update a mapping, check out the mapping and all mapping tasks that use it to ensure that changes to the mapping are propagated to the mapping tasks.

- Check in a reusable asset and all dependent assets in one operation.

  This ensures that changes to the asset and dependent assets are committed to the source control repository at the same time. It also ensures that users get the latest versions of the dependent assets when they pull the assets.

- Enter comments when you check in assets.

  When you check in assets, you might enter a release tag name in the **Summary** field and enter more descriptive comments in the **Description** field. When you do this, the **Git Summary** field in Informatica Intelligent Cloud Services shows the release tag that is associated with the asset.

- When you check in multiple assets at one time, limit the number of assets to 1000 or fewer.

  Checking in more than 1000 assets at one time can degrade performance between Informatica Intelligent Cloud Services and the GitHub repository service.

Undoing a checkout for another user

If you have the Admin role or your user role has the Force Undo Checkout feature privilege for the Administrator service, you can undo the checkout of an object that has been checked out by another user. You might need to undo the checkout of an object that has been checked out by another user if the user has checked out objects and goes on vacation or leaves the organization.

When you undo a checkout, the object reverts to the last version in the source control repository. The object's version history will not include a record of the checkout and undo checkout actions. If you think you might need the changed version later, make a copy of the object before you undo the checkout.

An undo action is not recursive. If you undo the checkout of a project or folder, the lock for the project or folder is released but the objects within the project or folder remain locked.

1. Open the service in which the user checked out the object.
2. On the **Explore** page, navigate to the object.

3. In the row that contains the object, click **Actions** and select **Undo Check Out**.

   The undo action releases the lock so that the object is available for checkout.

   Note: If an object was moved or renamed after it was checked out, undoing the checkout will restore the object's name and location to its name and location before it was checked out.

Rolling upgrades for Secure Agent services

Some Secure Agent Services support rolling upgrades. In a rolling upgrade, services that run on the agents within a Secure Agent group are upgraded sequentially. Therefore, while a service is being upgraded on one agent, the service remains available on other agents in the group.

The following Secure Agent service supports rolling upgrades:

- Process Server

Other services that run on the agents within a Secure Agent group are upgraded on each agent simultaneously. Therefore, they are unavailable while the agents in the group are being upgraded. They become available again when all agents in the group have been successfully upgraded.

Example

Your organization uses the following runtime environments:

- Secure Agent group A:

  Agent A1 runs Data Integration Server and Process Server.

  Agent A2 runs Data Integration Server, Mass Ingestion, and Process Server.

- Secure Agent group B:

  Agent B1 runs Data Integration Server, Mass Ingestion, and Process Server.

  Agent B2 runs Data Integration Server, Mass Ingestion, and Process Server.

When your organization is upgraded, Secure Agent groups A and B are upgraded simultaneously. Within each Secure Agent group, Process Server is upgraded sequentially. Therefore, while Process Server is being upgraded on agents A1 and B1, it remains up and running on agents A2 and B2. When the upgrade finishes on agents A1 and B1, Process Server is upgraded on agents A2 and B2.

Data Integration Server and Mass Ingestion do not support rolling upgrades. In groups A and B, Data Integration Server is upgraded on each agent simultaneously. In group B, Mass Ingestion is upgraded on agents B1 and B2 simultaneously.

Rolling upgrade error handling

If a service that supports rolling upgrades encounters an error during an upgrade, you can specify whether to continue or stop upgrading the service. Configure the error handling behavior on the **Settings** page.

You can select either of the following options:

In case of error, flag error and continue with upgrade

If an error occurs while a service is being upgraded, the service stops with an error on the agent in which it encountered the error. The upgrade then continues on the other agents within the group.

Warning: If you enable this option and the error occurs on all agents in the group, the service stops running on the Secure Agent group. This can cause job interruptions.

In case of error, stop upgrade

If an error occurs while a service is being upgraded, the service stops with an error on the agent in which it encountered the error. Upgrade of the service stops for all other agents in the group that have not already been upgraded. The agents that have not been upgraded continue to run the previous version of the service.

This is the default option.

To configure the error handling behavior, click **Edit** in the Upgrade Settings for Secure Agent Services area, select the appropriate option, and click **Save**.

Restart schedule configuration for Secure Agent services

Some Secure Agent services such as Process Server might need to be restarted after they are upgraded. You can configure a restart schedule for some of these services after a minor upgrade such as a monthly upgrade or a patch release. Configure the restart schedule on the **Settings** page.

When you configure a restart schedule, you select the day of the week and time in which to restart the services. For example, you might schedule the restarts for Sundays at 00:00 GMT.

You can configure a restart schedule for the following Secure Agent service:

- Process Server

To configure the schedule, click **Edit** in the Upgrade Settings for Secure Agent Services area, select a day and time, and click **Save**.

CHAPTER 8

Users and user groups

Configure users and user groups to allow access to your organization and assets.

A user is an individual account in Informatica Intelligent Cloud Services that allows secure access to an organization.

A user group is a group of user accounts in which all members of the group can perform the same tasks and have the same access rights for different types of assets.

Users and groups can perform tasks and access assets based on the roles that you assign to them. For more information about user roles, see Chapter 9, “User roles” on page 64.

Users

A user is an individual Informatica Intelligent Cloud Services account that allows secure access to an organization. A user can perform tasks and access assets based on the roles that are assigned to the user. You can assign roles directly to the user or to a group that the user is a member of.

Administrators can create and configure user accounts for the organization.

The **Users** page lists the users in your organization. To access the **Users** page, in Administrator, select **Users**.

The following image shows the **Users** page:

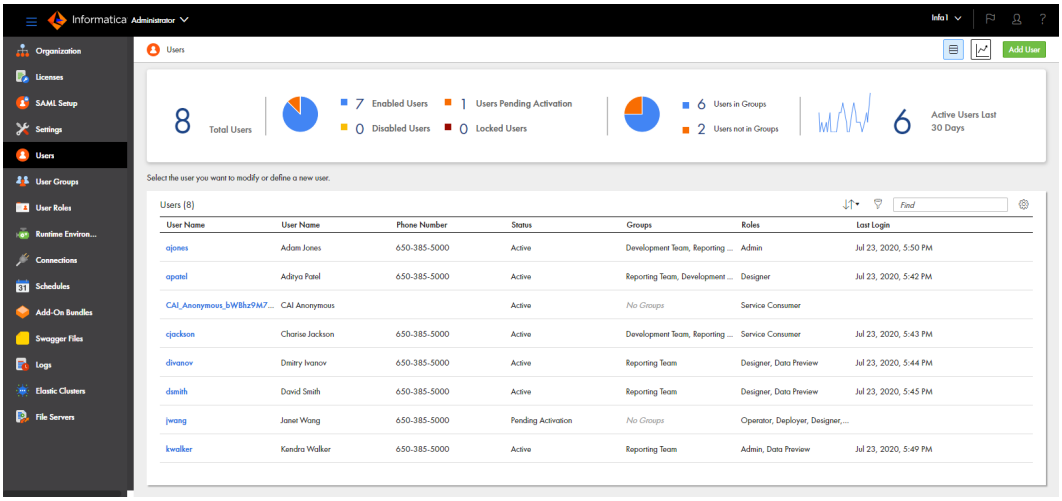

The **Users** page displays user statistics for the organization and lists each user. If you use Application Integration, the page also lists the Application Integration anonymous user and its status. To view detailed information about a user, click the user name.

You can perform the following tasks for a user:

- View and edit user details.
- Create a user.
- Assign and unassign services.
- Disable a user.
- Reset a user.
- Reassign a user's scheduled jobs to a different user.
- Delete a user.

User authentication

Informatica Intelligent Cloud Services uses different types of user authentication. Native users are authenticated through Informatica Intelligent Cloud Services. Salesforce, Microsoft Azure, and SAML users are authenticated through their identity providers.

Informatica Intelligent Cloud Services can use the following types of user authentication:

Native

Native users log in to Informatica Intelligent Cloud Services through the Informatica Intelligent Cloud Services login page using their user names and passwords. They are authenticated through Informatica Intelligent Cloud Services.

Salesforce

Salesforce users sign in to Informatica Intelligent Cloud Services through Salesforce or a Salesforce app. They are authenticated through Salesforce.

For more information about Salesforce authentication, see the help for the Salesforce connector in the Data Integration help.

Microsoft Azure

Microsoft Azure users sign in to Informatica Intelligent Cloud Services through Microsoft Azure. They are authenticated through Microsoft Azure.

For more information about Microsoft Azure authentication, see Chapter 4, “Ecosystem single sign-on” on page 28.

SAML

SAML users sign in to Informatica Intelligent Cloud Services through their identity provider. They are authenticated through their identity provider.

For more information about configuring SAML single sign-on, see Chapter 5, “SAML single sign-on” on page 30.

Application Integration anonymous user

If you have licensed Application Integration, Informatica Intelligent Cloud Services creates a system user called `CAI_Anonymous_<Organization_ID>`. Application Integration needs this user when you invoke an anonymous process that calls a Data Integration task.

Important: Do not edit or delete the Application Integration anonymous user if you need to invoke an anonymous process that calls a Data Integration task.

If you assign custom permissions to a Data Integration task and invoke the Data Integration task through an Application Integration process or a guide, you must complete either of the following tasks:

- Give the Application Integration anonymous user permission to run the associated Data Integration asset.
- Add the Application Integration anonymous user to a user group that has permission to run the associated Data Integration asset.

User statistics

If you have the Admin role or the "Read User" and "Audit log - view" privileges, you can view user statistics for your organization.

The statistics area on the **Users** page displays statistics such as the number of users in the organization, the number of users with each status, and the number of active users in a certain time period.

The following image shows the statistics area:

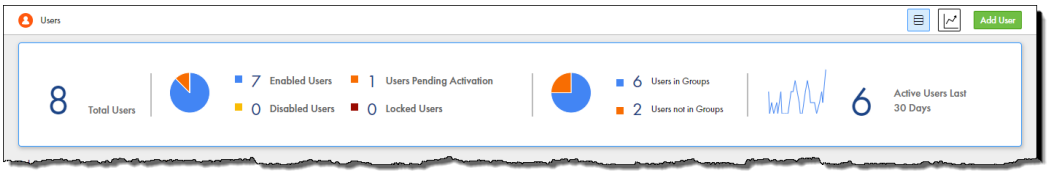

You can use the statistics area to filter the users on the **Users** page. For example, to display only users with the status Pending Activation, click **Users Pending Activation**. To list all users, click **Total Users**.

If you have the Admin role or the "Create User" and "Audit log - view" privileges, you can view a graph of the active users per day in the last 7, 30, or 90 days. To view the graph, click **Chart View** and select the appropriate time period. You can also download a report that lists the login date and time for each user during the time period.

To return to the list view of the **Users** page, click **List View**.

User details

You can configure user details such as user name, email, login settings, and assigned user groups and roles on the user details page. To display the user details page, in Administrator, select **Users**, and then click the user name.

The following image shows the user details page:

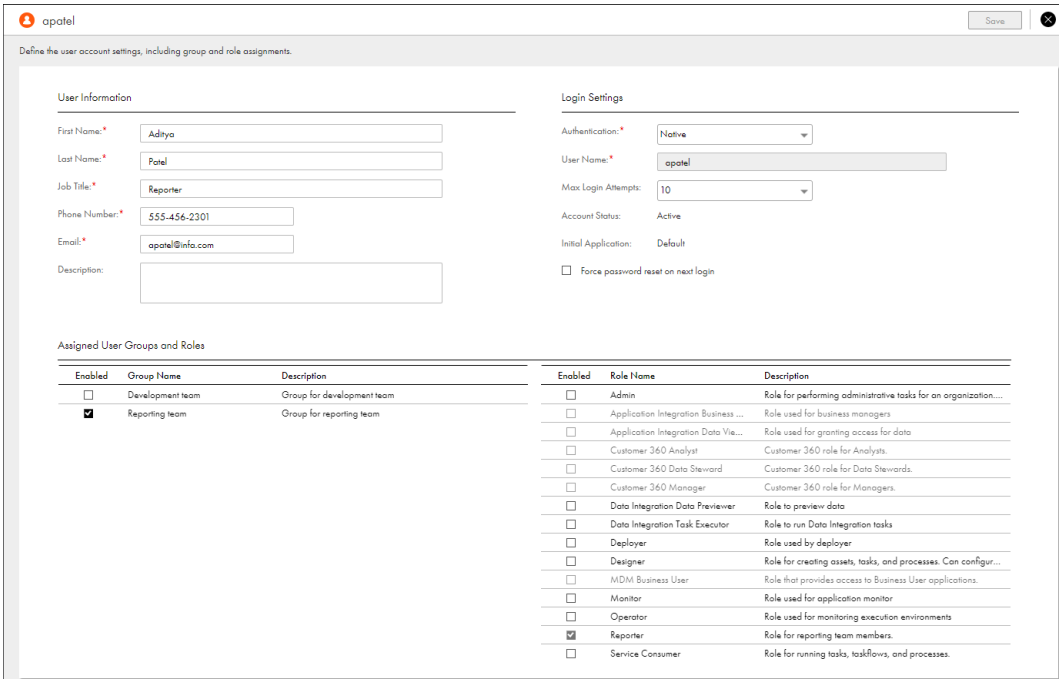

You can configure the following details for a user:

User information

The following table describes the user information:

Property	Description
First name	First or given name of the user.
Last name	Last or family name of the user.
Job title	User job title.
Phone number	Telephone number for the user.
Email	Email address of the user. Must be a valid email address in the format: `<local_part>@<domain>`. For example, jsmith@mycompany.com.
Description	Optional user description.

Login settings

The following table describes the login settings:

Property	Description
Authentication	Authentication method. Can be one of the following values: - Native. The user is authenticated through Informatica Intelligent Cloud Services. The user logs in through the Informatica Intelligent Cloud Services URL. - Salesforce. The user is authenticated through Salesforce and signs in through Salesforce or a Salesforce app. - Azure SSO. The user is authenticated and signs in through Microsoft Azure. - IDP with SAML. The user is authenticated and signs in through a SAML identity provider.
Activate using verification code / Activate using Salesforce OAuth	Account activation method for Salesforce users. Select one of the following options: - Activate using verification code. Select this option when the user signs in to Informatica Intelligent Cloud Services through a Salesforce app. When you select this option, the user receives an email with a verification code. The user account is activated when the user logs in to Salesforce, opens the Salesforce app, and enters the verification code. - Activate using Salesforce OAuth. Select this option to activate the user account using Salesforce OAuth. When you select this option, the user receives an email with a **Confirm Account** link. The user account is activated when the user clicks the **Confirm Account** link and enters the Salesforce user name and password. These options are displayed when the authentication method is Salesforce.
Environment	Salesforce organization environment, either production or sandbox. This option displayed when the user activation method is Salesforce OAuth.
User name	Informatica Intelligent Cloud Services user name. Must be unique within the Informatica Intelligent Cloud Services organization. You cannot change the name after you save the user. This property is displayed when the authentication method is Native.
Salesforce user name	Salesforce user name. Must be unique within the Informatica Intelligent Cloud Services organization. You cannot change the name after you save the user. For Salesforce users, the Informatica Intelligent Cloud Services user name is the same as the Salesforce user name unless that name is already used in the Informatica Intelligent Cloud Services organization. If the name is already used, then Informatica Intelligent Cloud Services appends the string ".Salesforce," ".Salesforce1," ".Salesforce2," etc. to the end of the Salesforce user name to form a unique Informatica Intelligent Cloud Services user name. This property is displayed when the authentication method is Salesforce.
Azure user name	Microsoft Azure user name. Must be unique within the Informatica Intelligent Cloud Services organization. You cannot change the name after you save the user. For Microsoft Azure users, the Informatica Intelligent Cloud Services user name is the same as the Azure user name unless that name is already used in the Informatica Intelligent Cloud Services organization. If the name is already used, then Informatica Intelligent Cloud Services appends the string ".Azure," ".Azure1," ".Azure2," etc. to the end of the Azure user name to form a unique Informatica Intelligent Cloud Services user name. This property is displayed when the authentication method is Azure SSO.

Property	Description
SAML user name	SAML user name. Must be unique within the Informatica Intelligent Cloud Services organization. You cannot change the name after you save the user. For SAML users, the Informatica Intelligent Cloud Services user name is the same as the SAML name identifier unless that name is already used in the Informatica Intelligent Cloud Services organization. If the name is already used, then Informatica Intelligent Cloud Services appends the string ".SAML," ".SAML1," ".SAML2," etc. to the end of the SAML name identifier to form a unique Informatica Intelligent Cloud Services user name. This property is displayed when the authentication method is IDP with SAML.
Max login attempts	Maximum number of login attempts that the user can make before the user is locked out. Select a number or "No Limit." If locked out, the user can click the **Forgot your password** link on the Login page, or the organization administrator can reset the user on the **Users** page. This property is displayed when the authentication method is Native.
Account status	Account status. Can be one of the following statuses: - Pending Activation. The user account has been created or reset, but the user has not yet activated the account. - Active. The user account has been created and validated, and the user can log in to Informatica Intelligent Cloud Services. - Locked. Applies to native user accounts. The account is locked because the user has exceeded the maximum number of login attempts. To unlock the user, the user can click the **Forgot your password** link on the Login page, or you can reset the user on the **Users** page. - Disabled. The user account has been disabled by an administrator. The user cannot log in to Informatica Intelligent Cloud Services.
Initial application	This field is reserved for future use.
Force password reset on next login	Forces the user to reset the password the next time the user tries to log in. This property is displayed when the authentication method is Native.

Assigned user groups and roles

You must assign at least one user group or role to each user. To assign or remove a user group or role, enable or disable the group or role, and then click **Save**.

When you assign a group to a user, all roles that are associated with the group become enabled. You cannot remove these roles individually. To remove the roles, you must remove the group.

Creating a user

Create a user on the **Users** page. When you create a user, the user status is set to Pending Activation or to Active based on the authentication method.

1. In Administrator, select **Users**.
2. Click **Add User**.
3. Enter the user information.
4. Enter the login settings:
   a. Select the authentication method.

b. For Salesforce users, specify whether to activate the user account using a verification code or Salesforce OAuth.

c. Enter the Informatica Intelligent Cloud Services user name or the user name in the third-party identity provider's system.

For native users, enter the Informatica Intelligent Cloud Services user name. For Salesforce, Microsoft Azure, or SAML users, enter the user name in the third-party identity provider's system.

The user name must be unique within the Informatica Intelligent Cloud Services organization. You cannot change the user name after you create a user.

d. For native users, select the maximum number of login attempts.

5. In the Assigned User Groups and Roles section, select the user groups and roles that you want to assign to the user.

You can assign system-defined and custom roles to a user. If you assign a group, the user inherits all roles that are associated with the group.

6. Click **Save**.

After you create a user, the user status is set as follows based on the authentication method:

- Native users are set to Pending Activation. The user receives an email to confirm the account. When the user clicks the **Confirm Account** link in the email, the user is prompted to set up a password and security question. When the user does this, the status changes to Active, and the user can log in to Informatica Intelligent Cloud Services.
- Salesforce users are set to Pending Activation.

  If you activate the user using a verification code, the user receives an email with a verification code. The user account is activated when the user logs in to Salesforce, opens the Salesforce app, and enters the verification code.

  If you activate the user using Salesforce OAuth, the user receives an email with a **Confirm Account** link. The user account is activated when the user clicks the **Confirm Account** link and enters the Salesforce user name and password.
- Microsoft Azure and SAML users are set to Active. The user can sign in through the user's identity provider.

Assigning and unassigning services

When you create a user, the user can access services based on the organization's licenses and the user's role. You can restrict the user's access to these services.

To allow or prevent a user from accessing certain services, you assign or unassign the services to the user. Assign and unassign services to a user on the **Users** page.

When you assign a service to a user, the service is visible on the **My Services** page. The user can access and use the service as long as the user's role allows this.

When you unassign a service, the user cannot see the service on the **My Services** page. The user cannot access or use the service regardless of the user's role.

For example, you want to allow an application developer with the Service Consumer role to use API Portal but not Data Integration or Application Integration. Assign the API Portal service to the user and unassign the Data Integration and Application Integration services. When you do this, the application developer can no longer see the Data Integration and Application Integration services on the **My Services** page. The application developer cannot use these services even though the Service Consumer role has privileges related to them.

1. In Administrator, select **Users**.

2. In the row that contains the user, click **Actions** and select **Assign Services**.
3. In the **Assign Services** dialog box, select the services that you want to assign to the user and deselect the services that you want to unassign.
4. Click **Save**.

Disabling a user

Disable a user on the **Users** page. When you disable a user, the user can no longer log in to Informatica Intelligent Cloud Services.

Before you disable a user, verify that the user did not schedule any tasks or taskflows. If you disable a user who has scheduled tasks or taskflows, the scheduled jobs fail.

When you disable a user, the user remains in the organization and in the Informatica Intelligent Cloud Services repository. You can view the user details, but you cannot edit them. Assets that the user created or updated also remain in the organization. On the Explore page, the Created by and Updated by columns indicate that the user is disabled.

1. In Administrator, select **Users**.
2. In the row that contains the user whom you want to disable, click **Actions** and select **Disable**.

Resetting a user

Reset a user on the **Users** page. You can reset a user whose account is disabled or locked. When you reset a user, the user status is set to Pending Activation or to Active based on the authentication method.

1. In Administrator, select **Users**.
2. In the row that contains the user, click **Actions** and select **Reset**.

After you reset a user, the user status is reset differently based on the authentication method:

- Native users are set to Pending Activation. The user receives an email to confirm the account. When the user clicks the **Confirm Account** link in the email, the user is prompted to reset the password and security question. The user can then log in to Informatica Intelligent Cloud Services.
- Salesforce users are set to Pending Activation.

  If you activated the user using a verification code, the user receives an email with a verification code. The user account is activated when the user logs in to Salesforce, opens the Salesforce app, and enters the verification code.

  If you activated the user using Salesforce OAuth, the user receives an email with a **Confirm Account** link. The user account is activated when the user clicks the **Confirm Account** link and enters the Salesforce user name and password.
- Microsoft Azure and SAML users are set to Active. The user can sign in through the user's identity provider.

Reassigning a user's scheduled jobs

Reassign a user's scheduled jobs on the **Users** page. You might want to reassign scheduled jobs when a user that has scheduled tasks or taskflows leaves the organization. You must reassign the user's scheduled jobs before you can delete the user.

The owner of a scheduled job is the last person that saves the scheduled task or taskflow. For example, in your organization, user Arun creates a schedule, user Beth creates a mapping task and assigns the schedule to the task, and then user Chandra updates and saves the task. Chandra becomes the owner of the scheduled

job. If Chandra leaves the organization, you must reassign her scheduled jobs to another user before you can delete her user account.

1. In Administrator, select **Users**.
2. In the row that contains the user, click **Actions** and select **Reassign Scheduled Jobs**.
3. Select a user to whom to reassign the scheduled jobs.

   The user you select must be an active user.
4. Click **Reassign**.

Deleting a user

Delete a user on the **Users** page. When you delete a user, the user is removed from the organization and from the Informatica Intelligent Cloud Services repository.

Before you can delete a user, you must reassign the user's scheduled jobs to a different user.

Note: You cannot reset a deleted user. If you think you might need to reactivate the user account, disable the user instead of deleting the user.

1. In Administrator, select **Users**.
2. In the row that contains the user whom you want to delete, click **Actions** and select **Delete**.
3. If the user is the owner of any scheduled tasks or taskflows, Administrator prompts you to reassign the jobs to a different user. Select the user to whom you want to reassign the jobs and click **Reassign and Delete**.

If the user did not own scheduled tasks or taskflows, Administrator deletes the user. If the user was the owner of any scheduled tasks or taskflows, Administrator reassigns the jobs and then deletes the user.

User groups

A user group is a group of users in which all members can perform the same tasks and have the same access rights for different types of assets. Members of a group can perform tasks and access assets based on the roles that you assign to the group.

Administrators can configure user groups for the organization.

The **User Groups** page displays a list of all user groups in the organization. To access the **User Groups** page, in Administrator, select **User Groups**.

The following image shows the **User Groups** page:

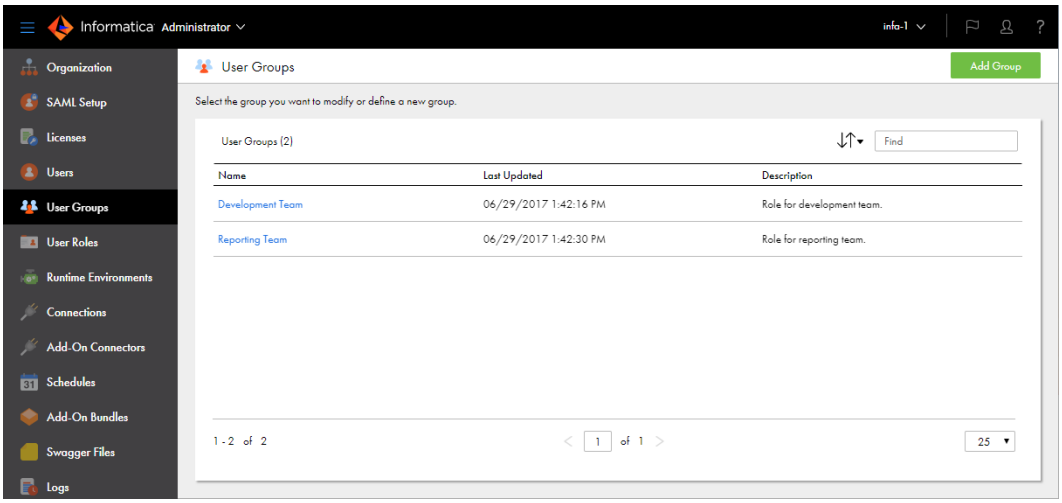

You can perform the following tasks for a user group:

- View and edit group details.
- Create a group.
- Rename a group.
- Delete a group.

User group details

You can configure details about a user group that include the group information, assigned roles, and group members on the group details page. To display the group details page, in Administrator, click **User Groups**, and then click the group name.

The following image shows the group details page:

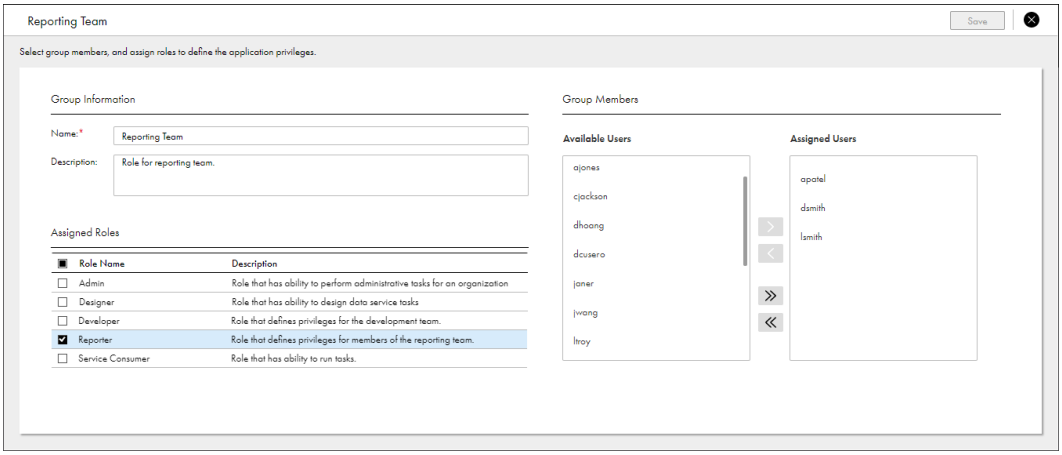

You can configure the following details for a user group:

Property	Description
Name	Required. Name of the user group. Must be unique within an organization. You can change the group name after you create it.
Description	Optional description for the user group.
Assigned roles	Roles that are assigned to all members of the group. You must assign at least one role to each group. To assign or remove a role, enable or disable the role, and then click **Save**.
Group members	Users who are assigned to the group. To assign a user to the group, move the user from the **Available Users** list to the **Assigned Users** list, and then click **Save**. To remove a user from the group, move the user from the **Assigned Users** list to the **Available Users** list, and then click **Save**. When you assign a user to a group, the user is automatically assigned all roles that are assigned to the group.

Creating a user group

Create a user group when multiple users in your organization need to perform the same tasks and need the same access rights for different types of assets. Group members can perform tasks and access assets based on the roles that you assign to the group. Create a user group on the **User Groups** page.

1. In Administrator, select **User Groups**.
2. Click **Add Group**.
3. Enter a group name and optional description.

   The group name must be unique within an organization.
4. In the Assigned Roles section, select the roles that you want to assign to the group.

   You can assign system-defined and custom roles to a group. The roles apply to all members of the group.
5. Optionally, assign users to the group.

   To assign a user to the group, move the user from the **Available Users** list to the **Assigned Users** list.

   You can also assign a user to a group when you create or edit a user.
6. Click **Save**.

Renaming a user group

Rename a user group on the **User Groups** page. You can also edit the user group and change the group name on the Group Details page.

1. In Administrator, select **User Groups**.
2. In the row that contains the user group, click **Actions** and select **Rename**.
3. Enter the new name and click **Save**.

Deleting a user group

Delete a user group on the **User Groups** page.

Tip: Before you delete a user group, verify that all group members have appropriate roles or are assigned to other groups so that they can continue to use Informatica Intelligent Cloud Services without interruption.

1. In Administrator, select **User Groups**.
2. In the row that contains the user group, click **Actions** and select **Delete**.

User configuration examples

The following examples illustrate ways in which you can configure users and user groups to control access to Informatica Intelligent Cloud Services according to your business needs.

For information about user roles, see Chapter 9, “User roles” on page 64.

You want your development team to create tasks and taskflows in Data Integration. The development team needs to view sample data in development, but you want to restrict access to production data.

1. Create a Developer role for the development team. Configure the role with all privileges for tasks and related assets, but only the Read privilege for connections.
2. Create a Development Team user group and add all members of the development team to the group.
3. Assign the Developer role to the Development Team group.
4. If possible, create development connections to sample data. If you have both development and production connections, configure the production connections so that the Development Team group does not have read permission for these connections. This prevents users in the Development Team group from using production connections in tasks.
5. After testing is complete and tasks are ready to move into production, have an administrator or other qualified user configure the tasks to use production connections.
6. Edit the Developer role and remove the privilege to run tasks. If development is complete for a task type, you can also remove the privileges to read and update the tasks. By removing the read privilege, you prevent users with the Developer role from accessing information about production tasks.

You have a reporting team that needs to run tasks in Data Integration, but does not have the technical knowledge to configure tasks safely.

1. Create a Reporter role for the reporting team. Configure the role with privileges to read and run tasks and taskflows, and privileges to read, create, and update schedules. Do not enable privileges to create, update, delete or set permissions on assets in the organization.
2. Create a Reporting Team user group and add all members of the reporting team to the group.
3. Assign the Reporter role to the Reporting Team group.

You want a security administrator who can assign roles and user groups and configure access control, but cannot create, edit, or run tasks.

1. Create a custom role called Security Administrator.
2. Edit the Security Administrator role and grant all privileges except the privileges to create, update, delete, and run tasks, connections, and schedules.

3. Assign the Security Administrator role to the security administrator.

You want to easily keep track of your organization administrators.

Create a user group called "Organization Administrators" and assign the Admin role to the group. Add all of your organization administrators to the group.

Your organization uses an OrderProcessing API to manage orders to a large supplier. This API consists of processes in Application Integration that include CreateOrder, ApproveOrder, and GetOrder. As an Admin, you want to restrict access to the ApproveOrder process to a few people.

1. Create a custom role called Approver. Configure the Run privilege for Application Integration Assets for the Approver role.
2. Create a user group called Order Approvers.
3. Assign the Approver role to the Order Approvers group.
4. Assign the Service Consumer role to the Order Approvers group. You must do this as the Service Consumer role can access and invoke processes.
5. Assign the users who need to be able to invoke ApproveOrder to the Order Approvers group.
6. In the Allowed Roles field of the ApproveOrder process, enter Approver.

Only members of the Order Approvers group will be able to invoke the ApproveOrder process.

You want an Application Integration developer to be able to perform all functions in the Application Integration Console except for viewing detailed process logs.

1. Create a role called Custom_Dev and configure the role with the following privileges:
   a. Select the Application Integration service, go to the **Assets** tab, and enable all CRUD privileges for **Application Integration Assets**.
   b. Go to the **Features** tab and add the Development, Console Administration, Publish Application Integration Assets, View Application Integration Console, and View Application Integration Designer privileges to the role.
   c. Select the Data Integration service, go to the **Assets** tab, and enable all CRUD privileges for the **Project** and **Folder** assets.
2. Assign the Custom_Dev role to the developer.

CHAPTER 9

User roles

A role is a collection of privileges that you can assign to users and groups. To ensure that every user can access assets and perform tasks in your organization, assign at least one role to each user or user group.

A role defines the privileges for different types of assets and service features. For example, users with the Designer role can create, read, update, delete, and set permissions on most types of data integration assets. However, they have no access to certain Administrator service features such as sub-organizations and audit logs.

Administrators can configure and assign roles for the organization.

You can assign the following types of roles:

System-defined

System-defined roles are pre-defined roles that define access privileges for the services that your organization uses. The system-defined roles that you can assign to users and groups vary based on your organization's licenses. You cannot edit or delete system-defined roles.

Custom roles

Custom roles are roles that you create to set privileges individually. To create custom roles, you need the appropriate license. You can edit and delete custom roles.

You can view both system-defined and custom roles on the **User Roles** page. The **User Roles** page displays a list of all roles in the organization. To access the **User Roles** page, in Administrator, select **User Roles**.

The following image shows the **User Roles** page:

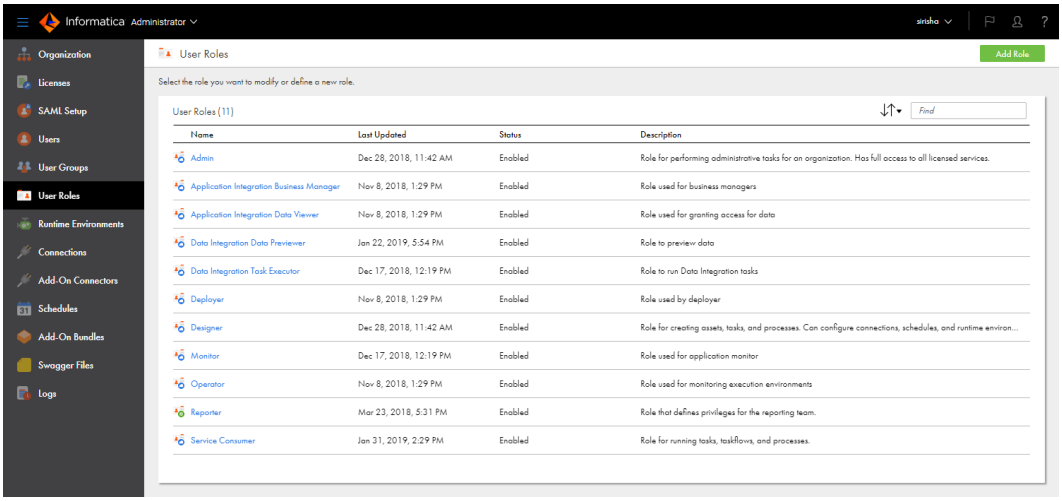

The Status column indicates whether the role is enabled or disabled for your organization. A role is disabled when the license expires.

You can assign multiple roles to a user or user group. When you assign multiple roles, the user or group inherits the access privileges associated with all of the roles.

Role details

The role details page displays information about a role, including the asset and feature privileges that are associated with the role. For system-defined roles, you can view the role information and privileges. For custom roles, you can view and change the role information and the assigned asset and feature privileges.

To display the role details page, in Administrator, select **User Roles**, and then click the role name.

The following image shows the role details page:

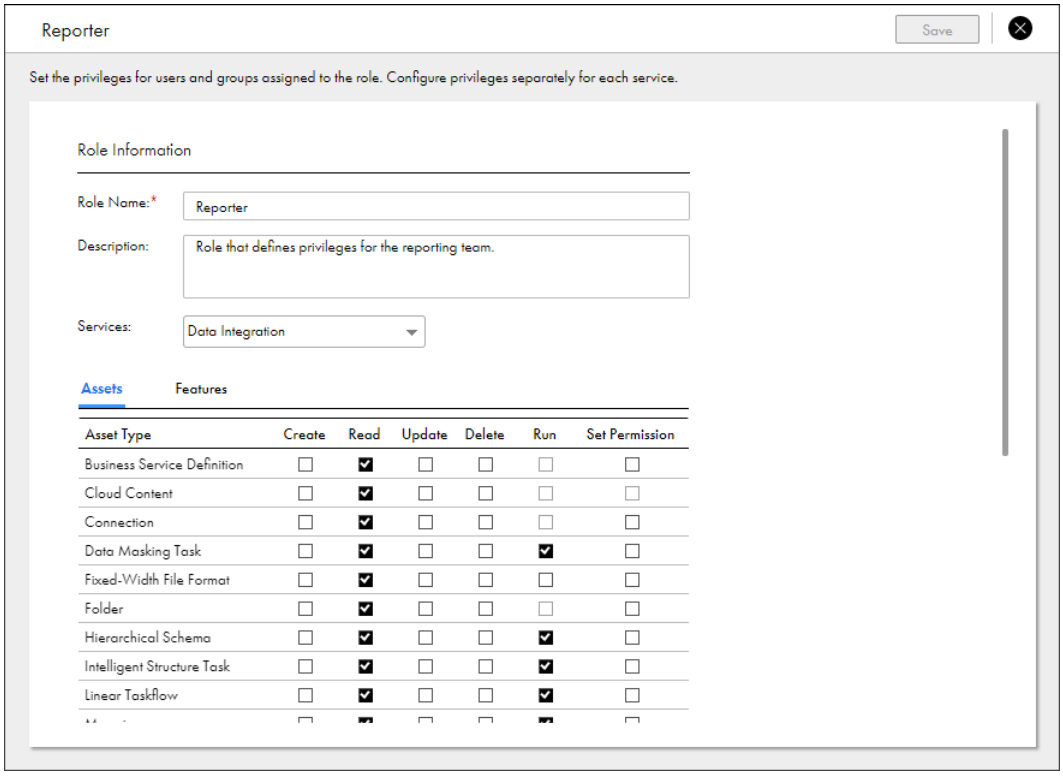

Each role has the following properties:

Role name

Name of the role. For custom roles, you can change the role name.

Description

Role description. For custom roles, you can change the role description.

Services

Name of the service for which privileges are enabled or disabled. Select a service to view the asset and feature privileges that are associated with the service.

If the license for a service expires, the service is marked as disabled. You can view the asset and feature privileges that are associated with a disabled service.

Assets

Asset privileges for the selected service. Asset privileges control access to different types of assets. For example, users with the Service Consumer role can view and run mappings in Data Integration, but they cannot create, update, delete, or set permissions on mappings.

The following table describes the asset privileges:

Privilege	Description
Create	Create assets of the selected type. For Secure Agents, this privilege allows users to download and install the Secure Agent. Requires the Read and Update privileges, which are automatically granted.
Read	Open assets of the selected type. For tasks, this privilege also allows users to use a connection or schedule in the task.
Update	Edit assets of the selected type. Requires the Read privilege, which is automatically granted.
Delete	Delete assets of the selected type.
Run	Run assets of the selected type. For the Data Integration service, users can run mappings, tasks, or taskflows. Users can also monitor, stop, and restart instances of the mapping, task, or taskflow. For the Hub Integration service, users can run publications or subscriptions.
Set permission	Configure permissions for assets of the selected type. For example, if you grant this privilege for projects, users with the role can select a project and enable other users and groups to read, update, delete, or change permissions for the selected project. To configure this privilege, your organization must have the appropriate license.

If a privilege does not apply to an asset type, the privilege is disabled. For example, the run privilege is disabled for folders.

For custom roles, you can enable and disable the asset privileges for a service as long as the service is not disabled.

Features

Feature privileges for the selected service. Feature privileges are general privileges that control the ability to use the features of a service. For example, users with the Designer role have the ability to perform data catalog discovery in Data Integration but not to preview data.

For custom roles, you can enable and disable feature privileges for a service as long as the service is not disabled.

Application Integration feature privileges

Use Application Integration feature privileges to create custom roles.

Important: You must assign the Folder and Project asset privileges to the user's role. To do this, select the Data Integration service and then select the CRUD options for the folder and project assets.

You can enable the following Application Integration feature privileges when you create a custom role:

Administration

Assign the Administration privilege to a role when you want the user to have complete design-time and run-time administrative access to the Application Integration and Application Integration Console.

Users with the Administration privilege can perform the following tasks:

- View, create, update, and delete all Application Integration assets.
- Manage and invoke services.
- Stop running processes.
- View instances and logs for deployed process.
- Deploy Process Developer BPR files to the Application Integration Console.
- Manage deployed catalogs.
- View WSDL files deployed across multiple systems.
- View Process Server metrics.

Note: The Application Integration Administration privilege does not give the user Informatica Intelligent Cloud Services-wide administrator privileges. For example, a user with the only the Application Integration Administration privilege will be unable to create sub-organizations.

Console Administration

Assign the Console Administration privilege to a role when you want the user to have near-complete access to the Application Integration Console.

Users with the Console Administration privilege can perform the following tasks:

- View instances for deployed process.
- Stop running processes.
- View deployed Process Developer BPRs and catalogs.
- View WSDL files deployed across multiple systems.
- View Process Server metrics.

Users with the Console Administration privilege cannot delploy BPR files.

Data Viewer

Assign the Data Viewer privilege to a user who needs to access detailed logs in the Application Integration Console.

For example, you could assign this privilege to a someone who needs to see all logs across the organization. You would not normally assign this role to a developer.

Note: The process logging level must be set to verbose to get detailed logs.

Development

Assign the Development privilege to developers who will occasionally need to debug processes.

Users with the Development privilege can perform the following tasks:

- View, create, update, and delete all Application Integration assets.
- Invoke services.
- View the Detailed Process Instance page on the Application Integration Console.
- Manage processes instances.

Monitoring

Assign the Monitoring privilege to a user who needs to view all parts of the Application Integration Console except for detailed logs.

Publish Application Integration Assets

Assign the Publish Application Integration Assets privilege to a user that needs to be able to publish Application Integration processes, guides, connections, or service connectors.

View Application Integration Console

Assign the View Application Integration Console privilege to a user who needs access to the Application Integration Console service. You must assign this privilege to any role that has privileges that include working on the Application Integration Console.

For example, you need to assign this privilege along with the Development privilege.

View Application Integration Designer

Assign the View Application Integration Designer privilege to a user who needs access to the Application Integration service. You must assign this privilege to any role that has privileges that include working on the Application Integration Console.

For example, you need to assign this privilege along with the Publish Application Integration Assets privilege.

Data Quality feature privileges

Use Data Quality feature privileges to grant users access to the preview functionality in data quality assets. You can enable the feature privileges when you create a custom role.

You can enable the following feature privileges for Data Quality:

Data Preview - Dictionaries

Enable the Data Preview - Dictionaries privilege on a role to enable a user to view the contents of a dictionary in the following cases:

- The user opens the dictionary from the Explore page.
- The user selects the dictionary in a Data Quality asset.

Data Preview - Test Panel

Enable the Data Preview - Test panel privilege on a role to enable a user to view data in the Test panel in a Data Quality asset.

The Data Quality feature privileges are enabled by default on the Administrator and Designer roles.

Note: The Data Preview - Dictionaries feature privilege and the Read privilege for dictionary assets work independently of each other. The Read privilege allows you to open the dictionary from the Explore page. The Data Preview - Dictionaries privilege allows you to view the dictionary data.

If you open a dictionary without the Data Preview - Dictionaries privilege, Data Quality displays a message to notify you that you do not have sufficient permissions to view the data.

System-defined roles

Informatica Intelligent Cloud Services provides system-defined roles that you can assign to users or user groups. You cannot change or delete the system-defined roles.

The system-defined roles that you can assign to users and groups vary based on your organization's licenses. For example, if your organization has no access to Application Integration or API Manager, you cannot assign the Deployer, Application Integration Business Manager, Application Integration Data Viewer, or Operator role to any user or group in your organization.

Assign system-defined roles to users and groups based on the tasks that they need to perform.

There are two types of system-defined roles:

- Cross-service roles define access privileges across multiple services.
- Service-specific roles define access privileges for one service or for a group of closely related services.

Cross-service roles

Cross-service roles are system-defined roles that define access privileges across multiple services.

For example, users with the Designer role can create assets and tasks in Data Integration, create assets in Cloud Integration Hub, create processes in Application Integration, and can also access the Application Integration Console. Users with the Monitor role can monitor Data Integration jobs, Cloud Integration Hub assets, and Application Integration process instances.

The following roles are cross-service roles:

- Admin
- Data Integration Data Previewer
- Deployer
- Designer
- Monitor
- Operator
- Service Consumer

The following table shows the services that each cross-service role can access:

	Admin role	Data Integration Data Previewer role*	Deployer role	Designer role	Monitor role	Operator role	Service Consumer role
Administrator	X	-	-	X	X	-	X
API Manager	X	-	X	-	-	-	X
API Portal	X	-	-	-	-	-	X
Application Integration	X	-	X	X	X	X	X

	Admin role	Data Integration Data Previewer role*	Deployer role	Designer role	Monitor role	Operator role	Service Consumer role
Application Integration Console	X	-	X	X	X	X	X
B2B Gateway	X	-	-	X	X	-	-
B2B Partners Portal	X	-	-	-	-	-	-
Data Integration	X	-	-	X	X	-	X
Data Quality	X	-	X	X	X	X	X
Data Profiling	X	-	-	X	X	X	-
Integration Hub	X	-	-	X	X	-	-
Monitor	X	-	-	X	X	-	-
Operational Insights	X	-	-	-	-	X	-

* The Data Integration Data Previewer role is a supplemental role that allows users to preview data in Data Integration and Data Profiling. It provides no access to services. Assign this role with another role that allows users to access Data Integration or Data Profiling.

In the preceding table, an "X" means that users with the role have access to the service. For example, users with the Admin role have access to all services.

Access privileges for cross-service roles

Assign cross-service roles to users who need access privileges for different services across Informatica Intelligent Cloud Services. Each cross-service role provides different access privileges.

Cross-service roles have the following access privileges:

Admin

Users with the Admin role have full access to all licensed services. They can perform all tasks in the organization when assigned both the Admin and Service Consumer roles.

The best practice is to assign the Admin role to one or two trusted users and assign the users to an administrative user group that has full permissions on all asset types. These users can act as alternative organization administrators and can help troubleshoot access control and other organization security issues.

Note: To provide full access to the API Manager service, including full privileges for OAuth 2.0 client management, assign the user both the Admin and Service Consumer roles.

Data Integration Data Previewer

Users with the Data Integration Data Previewer role can preview data when they select a source, target, or lookup object for use in a mapping or task in Data Integration. They can also view source object data when creating a profile or viewing profile results in Data Profiling.

The Data Integration Data Previewer role is a supplemental role. Assign this role with another role, such as the Designer role, to ensure that users can access Data Integration and Data Profiling.

Deployer

Users with the Deployer role can deploy Application Integration assets and manage APIs through API Manager. Assign this role in a production environment where deployment access is typically restricted.

Users with the Deployer privilege can view assets in Data Quality.

Note: To provide full access to the API Manager service, including full privileges for OAuth 2.0 client management, assign the user both the Deployer and Service Consumer roles.

The following table lists the services that users with the Deployer role can access and the access privileges associated with each service:

Service	Access Privileges
API Manager	Has full access to this service, including OAuth 2.0 client management privileges, when the Service Consumer role is also assigned.
Application Integration	Can view asset details.
Application Integration Console	Can deploy assets and view settings on the Processes, Logs, Server Configuration, Deployed Assets, and Resources pages. Can upload and deploy Process Developer-generated orchestration artifacts (BPRs).
Data Quality	Can view asset details.

Designer

Users with the Designer role can create assets, tasks, and processes. They can configure connections, schedules, and runtime environments. They can also monitor jobs and elastic clusters for the organization.

The following table lists the services that users with the Designer role can access and the access privileges associated with each service:

Service	Access Privileges
Administrator	Can configure connections, runtime environments, schedules, swagger files, and elastic configurations. Can install add-on connectors and install and uninstall add-on bundles. Can view upgrade settings for Secure Agent services. Can start and stop file servers, configure proxy servers, and view other file server settings.
Application Integration	Has full access to this service.
Application Integration Console	Can view and edit all settings except for server configuration properties.
B2B Gateway	Has full access to this service.

Service	Access Privileges
Data Integration	Has full access to this service.
Data Quality	Has full access to this service.
Data Profiling	Has full access to this service.
Integration Hub	Has full access to this service.
Monitor	Has full access to this service.

Monitor

Users with the Monitor role can monitor Data Integration jobs, Cloud Integration Hub assets, Data Quality assets, and Application Integration process instances for the organization.

The following table lists the services that users with the Monitor role can access and the access privileges associated with each service:

Service	Access Privileges
Administrator	Can view schedules and upgrade settings for Secure Agent services. Can start and stop file servers, configure proxy servers, and view other file server settings.
Application Integration	Can view asset details.
Application Integration Console	Can view settings.
B2B Gateway	Can view asset details.
Data Integration	Can view asset details.
Data Quality	Can view asset details.
Data Profiling	Can view asset details.
Integration Hub	Can view asset details.
Monitor	Can view data integration jobs and job details. Cannot view export or import jobs.

Operator

An Operator is responsible for process execution management and Process Server configuration updates. Users with the Operator role can view asset details but cannot modify them. They can manage process instances and modify some operational server parameters.

The following table lists the services that users with the Operator role can access and the access privileges associated with each service:

Service	Access Privileges
Application Integration	Can view asset details.
Application Integration Console	Can view and edit Process Server settings and some Cloud Server settings. For example, a user with the Operator role can create an alert service, but cannot view tenant details.
Data Quality	Can view asset details.
Data Profiling	Can view asset details.
Operational Insights	Can view cloud and domain infrastructure. Can edit domain and infrastructure Secure Agent alert settings. Can edit domain infrastructure, including registering domains.

Service Consumer

Users with the Service Consumer role can run tasks, taskflows, and processes but they cannot create or edit assets. Assign this role to users who need to execute Data Integration jobs and Application Integration processes through APIs.

Note: To provide full access to the API Manager service, assign the user both the Service Consumer and Deployer roles, or assign the user both the Service Consumer and Admin roles.

The following table lists the services that users with the Service Consumer role can access and the access privileges associated with each service:

Service	Access Privileges
Administrator	Can view schedules, swagger files, and upgrade settings for Secure Agent services. Can start and stop file servers, configure proxy servers, and view other file server settings.
API Manager	Has full access to this service when the Deployer or the Admin role is also assigned.
API Portal	Has full access to this service.
Application Integration	Can invoke Application Integration processes.
Data Integration	Can view tasks, run tasks, test-run mappings, run taskflows, and download workflow XML.
Data Quality	Can view asset details.

Service-specific roles

Service-specific roles are system-defined roles that define access privileges for one service or for a group of closely related services. For example, the service-specific roles for Application Integration provide access to both Application Integration and Application Integration Console.

Assign service-specific roles to users who do not need access across multiple services. Service-specific roles have different access privileges based on the services to which they apply.

The following table lists the service-specific roles for each service that uses them:

Service	Service-Specific Roles
Application Integration	Application Integration Business Manager Application Integration Data Viewer
Data Integration	Data Integration Task Executor
Reference 360	Reference 360 Administrator Reference 360 Business Analyst Reference 360 Business Steward Reference 360 Planner Reference 360 Primary Owner Reference 360 Stakeholder
Customer 360	Customer 360 Analyst Customer 360 Manager Customer 360 Data Steward MDM Business User
Business 360 Console	MDM Designer

Access privileges for Application Integration roles

Assign Application Integration roles to users who need access privileges for Application Integration and Application Integration Console. Each role provides different access privileges.

The following service-specific roles define access privileges for Application Integration and Application Integration Console:

Application Integration Business Manager

An Application Integration Business Manager monitors business activity. Users with the Application Integration Business Manager role can view information about assets and process instances, but they cannot change them.

The following table lists the services that users with the Application Integration Business Manager role can access and the access privileges associated with each service:

Service	Access Privileges
Application Integration	Can view folder and asset lists and asset details.
Application Integration Console	Can access the Processes page.

Application Integration Data Viewer

Users with the Application Integration Data Viewer role can view detailed logs in the Application Integration Console service.

Note: The logging level of an artifact must be set to verbose for a user to view detailed logs.

The Application Integration Data Viewer role is a supplemental role. Assign this role along with at least one other role. For example, if you want a user with the Designer role to view detailed Process Server

logs, assign the user the Application Integration Data Viewer and the Designer roles, and set the Process Server logging level to verbose.

Access privileges for Data Integration roles

The Data Integration Task Executor role defines access privileges for Data Integration. Users with the Data Integration Task Executor role can run tasks and taskflows and test-run mappings in Data Integration. They can also monitor data integration jobs.

The following table lists the services that users with the Data Integration Task Executor role can access and the access privileges associated with each service:

Service	Access Privileges
Administrator	Can view schedules and upgrade settings for Secure Agent services. Can start and stop file servers, configure proxy servers, and view other file server settings.
Data Integration	Can view assets and asset details, run tasks and taskflows, and test-run mappings. Can view user's own data integration jobs and job details, start and stop user's own jobs, and download session logs. Cannot view export or import jobs.
Monitor	Can view data integration jobs and job details, start and stop data integration jobs, and download session logs. Cannot view export or import jobs.

Access privileges for Reference 360 roles

Assign Reference 360 roles to users who need access privileges for Reference 360. Each role provides different access privileges.

The following service-specific roles define access privileges for Reference 360:

Reference 360 Administrator

Users with the Reference 360 Administrator role configure the Reference 360 environment.

Reference 360 Business Analyst

Users with the Reference 360 Business Analyst role view and analyze Reference 360 assets. They cannot propose changes to assets.

Reference 360 Business Steward

Users with the Reference 360 Business Steward role are subject matter experts for reference data. They create and manage code values in code lists and value mappings in crosswalks. They are responsible for approving changes proposed by other users. They can send their own changes for approval or directly publish their changes without approval. They can assign users access to crosswalks.

Reference 360 Planner

Users with the Reference 360 Planner role create and manage hierarchies. They assign users access to hierarchies.

Reference 360 Primary Owner

Users with the Reference 360 Primary Owner role create and define reference data structures, such as reference data sets and code lists. They can delete code lists and propose changes to code values in code lists. The user with the Business Steward role must approve the proposed changes. Primary owners can also assign users access to code lists and reference data sets.

Reference 360 Stakeholder

Users with the Reference 360 Stakeholder role propose changes to code values. The user with the Business Steward role must approve the proposed changes.

For more information about these roles, see the Reference 360 help.

Access privileges for Customer 360 roles

Assign Customer 360 roles to the users who need access privileges for Customer 360. Each role provides different access privileges.

The following service-specific roles define access privileges for Customer 360:

Customer 360 Analyst

Users with the Customer 360 Analyst role can create and edit records in Customer 360. When a Customer 360 Analyst creates or edits a record, the changes trigger a review process that requires approval from a Customer 360 Manager.

Customer 360 Manager

Users with the Customer 360 Manager role can review and approve customer records or update customer records. They can also create or edit records without approval.

Customer 360 Data Steward

Users with the Customer 360 Data Steward role can perform any task in Customer 360. They can create and edit records without approval, run jobs, and review and approve customer records.

MDM Business User

Users with the MDM Business User role can view records in Customer 360. They cannot create or edit records in Customer 360.

Access privileges for Business 360 Console roles

Assign Business 360 Console roles to the users who need access privileges for Business 360 Console.

The following service-specific role defines access privileges for Business 360 Console:

MDM Designer

Users with the MDM Designer role can define reference data in Business 360 Console.

Custom roles

A custom role is a role that you create based on the needs of your organization. For example, you might want to create a custom administrative role that can configure roles, user groups, and access control, but cannot create, edit, or run data integration tasks.

To create custom roles, your organization must have the appropriate license.

You can edit and delete custom roles after you create them.

You might want to edit custom roles when your organization gets a new license. Edit the roles to grant access to new asset types and features. Informatica Intelligent Cloud Services does not grant additional privileges to custom roles when your organization gets a new license.

Creating a custom role

Create a custom role on the **User Roles** page. When you create a role, you configure the privileges that are associated with the role. You configure privileges separately for each service.

1. In Administrator, select **User Roles**.
2. Click **Add Role**.
3. Enter a role name and optional description.
4. In the **Services** field, select the service for which you want to configure privileges.

   For example, to configure privileges for Data Integration, select **Data Integration**. To configure administrative privileges, select **Administrator**.
5. To configure the asset privileges, select **Assets**, and enable the appropriate privileges for each asset type.

   For example, to enable users with the role to create folders, enable **Create** next to **Folder**.

   To revoke an asset privilege, disable the privilege.
6. To configure the feature privileges, select **Features**, and enable the appropriate privileges.

   For example, to enable users with the role to import assets, enable **Asset - import**.

   To revoke a feature privilege, disable the privilege.
7. Repeat steps 4 through 6 for each service.
8. Click **Save**.

After you create a role, you can assign it to a user or user group. To assign the role to a user or group, edit the user or group.

Deleting a custom role

Delete a custom role on the **User Roles** page. You cannot delete a custom role if it is assigned to any user or user group. You cannot delete a system-defined role.

1. In Administrator, select **User Roles**.
2. In the row that contains the role that you want to delete, click **Actions** and select **Delete**.

B2B Partners Portal user roles

If your organization uses B2B Gateway, you might want to enable access to B2B Partners Portal for your external trading partners. To give your trading partners access to B2B Partners Portal, create a custom role and assign it to partner users.

When you create a custom role for partner users, name the role so that you know it is for B2B Partners Portal users. For example, you might name the role "B2B Partners Portal User."

Optionally, you can give the role a description. Clearly describe the role so that you know it is for users from partner companies. For example, you might describe the role as "Role for users from partner companies to access the B2B Partners Portal service."

When you create a custom role for B2B Partners Portal users, enable the Partners Portal feature privilege for the B2B Partners Portal service. For more information about creating custom roles, see "Creating a custom role" on page 77.

Assign the custom role to users from your partner companies. You only need to create one role for B2B Partners Portal users. Assign the same role to all external B2B Partners Portal users.

CHAPTER 10

Permissions

Permissions determine the access rights that a user has for a Secure Agent, Secure Agent group, connection, schedule, or asset. Permissions add additional or custom security for an object. Permissions define which users and groups can read, update, delete, execute, and change permissions on the object.

To configure permissions on an object, you need the following licenses and privileges:

- To configure permissions at the project level for all assets in a project, your organization must have the Set/Unset Security Permissions at Project Level license.
- To configure permissions at the folder level for all assets in a folder, your organization must have the Set/Unset Security Permissions at Folder Level license.
- To configure permissions on individual assets, your organization must have the Fine Grained Security license.
- The role assigned to your user account or to a group in which you are a member must have the Set Permission privilege for the object type. For example, to configure permissions on a Secure Agent, you must be assigned a role that has the Set Permission privilege for Secure Agents.

To configure permissions on an object, navigate to the object and set the appropriate permissions. For example, you want only users in the Development Team user group to have access to assets in the Development Data folder. Navigate to the folder, edit the permissions, and grant the Development Team user group permissions on the folder.

Permissions apply to the objects for which you configure them but not to copies of the object. Therefore, when you copy or export an asset, the permissions are not copied or exported with the asset. For example, you export a mapping task in which only user rjones has execute permission. When you import the mapping task, the imported mapping has no permissions assigned to it. Therefore, any user with privileges to run mapping tasks can run the imported task.

You can configure the following permissions on an object:

Permission	Description
Read	Open and view the object. If the object is source controlled, this permission allows the user or group to pull or check out the object from the source control repository. If you select a task, this permission also allows the user or group to use a connection or schedule in the task.
Update	Edit the object. If the object is source controlled, this permission allows the user or group to check in, check out, pull, unlink, or roll back the object. Requires read permission, which is automatically granted.

Permission	Description
Delete	Delete the object.
Execute	Run the object. Applies to mappings, tasks, and taskflows. Monitor, stop, and restart instances of the mapping, task, or taskflow.
Change permissions	Change the permissions that are assigned to the object.

Note: These permissions control permissions within Informatica Intelligent Cloud Services. They do not control operating system permissions, such as the ability to start, stop, or configure the Secure Agent on Windows or Linux.

Rules and guidelines for permissions

Use the following rules and guidelines for permissions:

- When you configure permissions on an object, verify that the user or group to which you grant permissions is assigned a role with the appropriate privileges for the object type. For example, if you grant a user with the Service Consumer role Update privilege on a particular folder, the user cannot update the folder because the Service Consumer role does not have update privileges for folders.
- To edit an asset, the user must have read permission on all assets used within the asset. For example, when you assign a user Read and Update permissions on a synchronization task, verify that the user also has Read permission on the connections, mapplets, schedules, and saved queries that are used in the task.
- When a user edits a task, assets without Read permission are not displayed. To avoid unexpected results, the user should cancel all changes and avoid editing the task until the user is granted the appropriate Read permissions.
- When configuring a taskflow, a user needs Execute permission on all tasks to be added to the taskflow.
- To edit a taskflow, a user needs Execute permission on all tasks in the taskflow. Without Execute permission on all tasks, the user cannot save changes to the taskflow.
- To run a taskflow, a user needs Read and Execute permissions on taskflows.
- To monitor jobs or to stop a running job, a user needs Execute permission on the mapping, task, or taskflow.
- If you assign custom permissions to a Data Integration task and invoke the Data Integration task through an Application Integration process or a guide, you must complete either of the following tasks:
  - Give the Application Integration anonymous user permission to run the associated Data Integration asset.
  - Add the Application Integration anonymous user to a user group that has permission to run the associated Data Integration asset.

Configuring permissions

You can configure permissions on an object if you are assigned a role with the Set Permission privilege for the object type. For example, to configure permissions on a folder, you must be assigned a role that has the Set Permission privilege for folders.

1. Navigate to the object for which you want to configure permissions.

   For example:

   - To configure permissions on a Secure Agent or Secure Agent group, in Administrator, select **Runtime Environments**.
   - To configure permissions on a connection, in Administrator, select **Connections**.
   - To configure permissions on a mapping, in Data Integration, open the project and folder that contain the mapping.

2. In the row that contains the object, either click **Actions** and select **Permissions**, or click the **Change Permission** icon.

   The **Permissions** dialog box lists the users and groups that have permissions on the object.

   If the **Permissions** dialog box lists no users or groups, then no permissions are configured for the object. Any user with appropriate privileges for the object type can access the object.

   The following image shows the **Permissions** dialog box for a mapping:

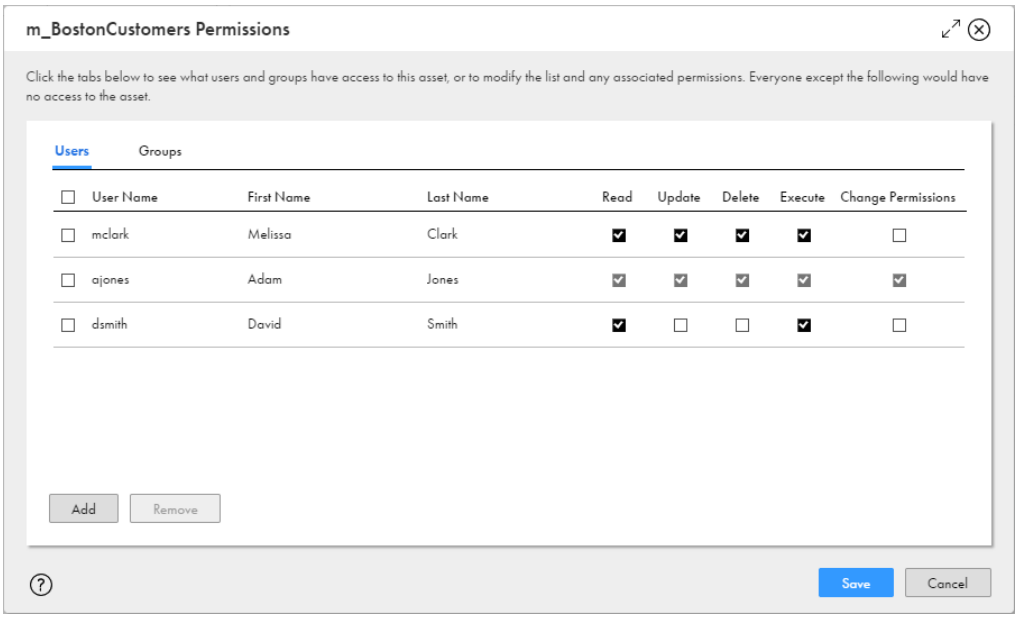

3. To configure user permissions on the object:

   a. Select **Users**.

   b. If the user does not appear in the **Users** list, click **Add**, and select a user.

   c. Enable or disable the appropriate permissions on the user.

   Note: When you grant any user permissions on the object, Informatica Intelligent Cloud Services also adds you as a user with permissions on the object. This prevents you from losing access to the object when you configure permissions.

4. To configure user group permissions on the object:
   a. Select **Groups**.
   b. If the group does not appear in the **Groups** list, click **Add**, and select a group.
   c. Enable or disable the appropriate permissions on the group.

   Note: When you grant any group permissions on the object, Informatica Intelligent Cloud Services also adds you as a user with permissions on the object. This prevents you from losing access to the object when you configure permissions.
5. To remove all permissions restrictions for the object, remove all users and groups from the **Permissions** dialog box.

   When you remove all users and groups, any user with appropriate privileges for the object type can access the object.
6. Click **Save**.

CHAPTER 11

Runtime environments

A runtime environment is the execution platform that runs a data integration or application integration task. You must have at least one runtime environment in each organization so that users in the organization can run tasks.

A runtime environment consists of one or more Secure Agents. A Secure Agent is a lightweight program that runs all tasks and enables secure communication across the firewall between your organization and Informatica Intelligent Cloud Services.

You can set up runtime environments in the following ways:

License the Informatica Cloud Hosted Agent.

If you license the Hosted Agent, you run tasks within the Informatica Cloud hosting facility. Informatica maintains the Hosted Agent runtime environment and agents.

Create one or more Secure Agent groups.

You can download and install one or more Secure Agents to run within your network or in a cloud computing services environment such as Amazon Web Services or Microsoft Azure. You can install one Secure Agent on each physical or virtual machine.

When you install a Secure Agent, it is added to its own group by default. If you have the Secure Agent Cluster license, you can add multiple agents to one Secure Agent group.

When you configure a connection or some types of tasks, you specify the runtime environment to use. The runtime environment determines which agent runs the tasks at run time. If the runtime environment is the Hosted Agent, the Hosted Agent runs the tasks. If the runtime environment is a Secure Agent group, any available agent in the group can run the tasks.

To use the runtime environment to run a mapping task that is based on an elastic mapping, the runtime environment must be associated with an elastic configuration. A Secure Agent uses the configuration to push data processing to an elastic cluster.

Hosted Agent

If your organization has the Cloud Runtime license, you can use the Hosted Agent to run tasks. The Hosted Agent can run synchronization and mapping tasks that use certain connectors.

You cannot use the Hosted Agent to run a mapping task that is based on an elastic mapping.

Informatica Cloud Data Integration manages the Hosted Agent runtime environment, so you cannot add, delete, or configure a Hosted Agent.

The Hosted Agent can run synchronization, mapping, and replication tasks that use certain connectors:

- Amazon Athena Connector
- Amazon Aurora Connector
- Amazon DynamoDB Connector
- Amazon Redshift Connector
- Amazon Redshift V2 Connector
- Amazon S3 Connector
- Amazon S3 V2 Connector
- Box Connector
- Box Oauth Connector
- CDM Folders Connector
- Cloud Integration Hub
- Concur V2 Connector
- Coupa V2 Connector
- DB2 Warehouse on Cloud Connector
- Eloqua Bulk API Connector
- Google Analytics Connector
- Google Big Query Connector
- Google Big Query V2 Connector
- Google Cloud Spanner Connector
- Google Cloud Storage Connector
- Google Cloud Storage V2 Connector
- Marketo V3 Connector
- Microsoft Azure Blob Connector
- Microsoft Azure Blob Storage V2 Connector
- Microsoft Azure Blob Storage V3 Connector
- Microsoft Azure Cosmos DB SQL API Connector
- Microsoft Azure Data Lake Connector
- Microsoft Azure Data Lake Store Gen2 Connector
- Microsoft Azure Data Lake Store V2 Connector
- Microsoft Azure Data Lake Store V3 Connector
- Microsoft Azure Data Warehouse Connector
- Microsoft Azure SQL Data Warehouse V2 Connector
- Microsoft Azure SQL Data Warehouse V3 Connector
- Microsoft Dynamics 365 for Operations Connector
- Microsoft Dynamics 365 for Sales Connector
- Microsoft SQL Server Connector
- Mock Connector
- MongoDB Connector

- MySQL Connector
- NetSuite Connector
- NetSuite V2 Connector
- Oracle Connector
- PostgreSQL Connector
- REST V2 Connector
- Salesforce Connector
- Salesforce Marketing Cloud Connector
- Salesforce Oauth Connector
- ServiceNow Connector
- Snowflake Cloud Data Warehouse V2 Connector
- Snowflake Connector
- SuccessFactors ODATA Connector
- SuccessFactors SOAP Connector
- SugarCRM REST Connector
- UltiPro Connector
- Workday V2 Connector
- Zendesk Connector
- Zendesk V2 Connector

Note: The Hosted Agent support is specific to connectors. For more information, see the help for the relevant connector.

Secure Agent groups

Use a Secure Agent group as the runtime environment when you need to access data on-premises or when you want to access data in a cloud computing services environment without using the Hosted Agent. When you select a Secure Agent group as the runtime environment for a connection or task, a Secure Agent within the group runs the tasks.

Create Secure Agent groups to accomplish the following goals:

Prevent the activities of one department from affecting another department.

To prevent the activities of one department from impacting a different department, create separate Secure Agent groups for each department. For example, users in the sales department run 10 times as many tasks as users in the finance department, but the finance tasks are more time critical. To prevent the sales tasks from impacting the finance tasks, create separate Secure Agent groups for each department. Then assign the sales tasks to one runtime environment and the finance tasks to the other runtime environment.

Separate tasks by environment.

You can create different Secure Agent groups for test and production environments. When you configure a connection, you can associate it with the test or production database by choosing the appropriate Secure Agent group as the runtime environment.

When you create a Secure Agent group, all users in the organization can select the Secure Agent group as the runtime environment.

You can add and remove Secure Agents from a group. Based on your license, you can also perform the following actions:

- If you have the Secure Agent Cluster license, you can add multiple agents to a Secure Agent group.
- If you have the Organization Hierarchy license, you can share a Secure Agent group with your sub-organizations.

Note: If you use the runtime environment to run a mapping task that is based on an elastic mapping, the Secure Agent group must have only one Secure Agent.

If you need to access output files on the Secure Agent machine, you can view the **All Jobs** page in Monitor or the **My Jobs** page in Data Integration to determine where a task ran.

Secure Agent groups with multiple agents

When you create a Secure Agent, it is added to its own group by default. If you have the Secure Agent Cluster license, you can add multiple agents to one Secure Agent group. All agents within a group must be of the same type, for example, all agents that run within your network or all agents that run on Amazon EC2 machines.

Add multiple agents to a group to achieve the following goals:

Balance the workload across machines.

Add multiple agents to a group to balance the distribution of tasks across machines. When the runtime environment is a Secure Agent group with multiple agents, the group dispatches tasks to the available agents in a round-robin fashion.

Improve scalability for connections and tasks.

When you create a connection or task, you select the runtime environment to use. If the runtime environment is a Secure Agent group with multiple agents, the tasks can run if any Secure Agent in the group is up and running. You do not need to change connection or task properties when you add or remove an agent or if an agent in the group stops running.

When you add multiple agents to a group, ensure that all of the Secure Agents are of the same type. For example, your organization installs four Secure Agents on physical machines within your network and two Secure Agents on Amazon EC2 machines. You can create a Secure Agent group that contains some or all of the local agents and a different group that contains the EC2 agents. Do not create a group that contains both a local agent and an EC2 agent.

If you need to access output files on the Secure Agent machine, you can view the job details to determine which Secure Agent ran the task. To view job details, open Monitor, select **All Jobs**, and click the job name.

Service assignment for Secure Agent groups

By default, when you create a Secure Agent group, all services that your organization uses can use the group. If your organization uses multiple services, the demand on the Secure Agent group can be high. To reduce the potential demand on a Secure Agent group, you can enable and disable specific Secure Agent services for the group.

The services that you enable and disable for a Secure Agent group are the Secure Agent services, which are different from the Informatica Intelligent Cloud Services. For example, if you want to use the agents in a group only for Operational Insights, enable the OI Data Collector service for the group and disable all other services. For more information about Secure Agent services, see Chapter 13, “Secure Agent services” on page 114.

You can perform the following actions:

Enable services for a Secure Agent group.

Enable services when you want the agents in the group to run the connections, tasks, processes, or product features associated with a service or set of services. When you enable a service, the service starts on each agent in the Secure Agent group.

Disable services for a Secure Agent group.

Disable services when you do not want the agents in the group to run the connections, tasks, processes, or product features associated with a service or set of services. When you disable a service, the service stops on each agent in the Secure Agent group. Any connection, task, process, or product feature that uses the Secure Agent group as the runtime environment no longer runs.

Enable or disable services for a Secure Agent group on the **Runtime Environments** page.

The following image shows the **Runtime Environments** page:

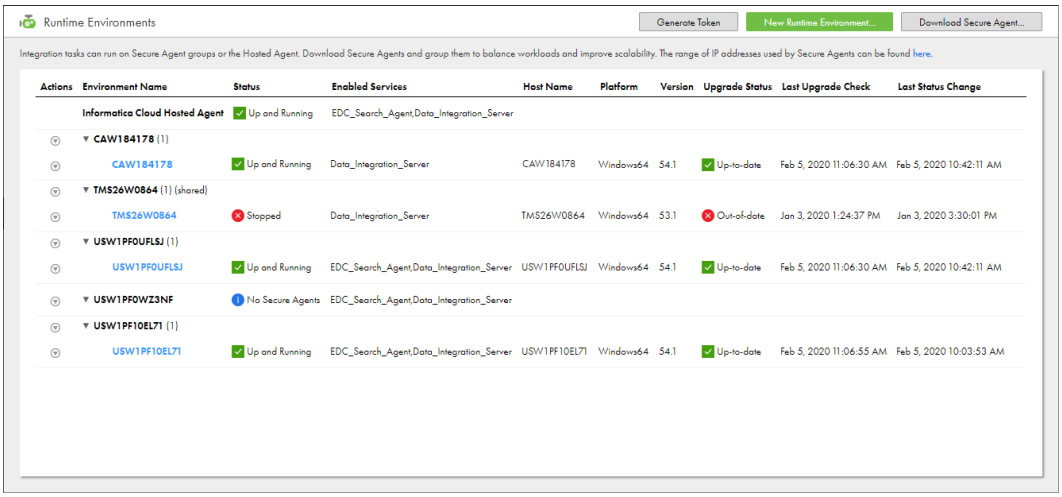

The Enabled Services column indicates which services are enabled for the Secure Agent group. The Enabled Services column for the Hosted Agent lists all Secure Agent services that your organization is licensed to use. To enable or disable a service, expand the **Actions** menu for the Secure Agent group, and select **Enable or Disable Services**.

After you make service assignments for a Secure Agent group, you might add or remove agents. When you add a Secure Agent to a group, the agent inherits the service assignments of the group that you add it to.

Example

Your organization uses Data Integration and has licenses for mass ingestion and for Enterprise Data Catalog data discovery. The organization uses the following Secure Agent groups:

- Group 1: Secure Agent 1, Secure Agent 2 , Secure Agent 3
- Group 2: Secure Agent 4
- Group 3: Secure Agent 5

By default, users in your organization can select any group as the runtime environment for any connection or any task, including file ingestion tasks. An administrator can also select any group as the runtime environment for integration with Enterprise Data Catalog.

To balance the load across Secure Agent groups, you want might want to reserve Group 1 for Data Integration tasks except file ingestion tasks, Group 2 for file ingestion tasks, and Group 3 for data catalog discovery.

Therefore, you enable and disable the following Secure Agent services:

Secure Agent Group	Enabled Services	Disabled Services
Group 1	Data Integration Server	Mass Ingestion, EDC Search Agent
Group 2	Mass Ingestion	Data Integration Server, EDC Search Agent
Group 3	EDC Search Agent	Data Integration Server, Mass Ingestion

To avoid task and feature failures, you must also verify the following settings:

- All Data Integration tasks except file ingestion tasks use Group 1 as the runtime environment. All connections that these tasks use also use Group 1 as the runtime environment.
- All file ingestion tasks use Group 2 as the runtime environment. All connections that these tasks use also use Group 2 as the runtime environment.
- On the **Organization** page in Administrator, the Enterprise Data Catalog integration properties use Group 3 as the runtime environment.

Service assignment guidelines

Use the following guidelines when you enable and disable services for a Secure Agent group:

- Before you disable a service, verify that no connection, task, or process that uses the group as the runtime environment requires the service.

  If a connection, task, or process has a Secure Agent group selected as the runtime environment and you disable a required service, the task or process cannot run. For example, the connection for a mapping source uses runtime environment RuntimeEnv1. If you disable Data Integration Server on RuntimeEnv1, the mapping task fails at run time.
- Before you disable a service, verify that no feature that uses the group as the runtime environment requires the service.

  If a feature has a Secure Agent group selected as the runtime environment and you disable a required service, the feature cannot be used. For example, the runtime environment for Enterprise Data Catalog integration is set to RuntimeEnv2. If you disable EDC Search Service on RuntimeEnv2, you can no longer perform data catalog discovery.
- When you create a connection, select a runtime environment in which the required services are enabled.

  For example, you want to create an Advanced SFTP connection for a file ingestion task target. When you create the connection, select a runtime environment in which the Mass Ingestion service is enabled.
- Do not disable a service to temporarily stop the service on a Secure Agent. For information about temporarily stopping a service on a Secure Agent, see "Stopping and starting services on a Secure Agent" on page 96.

Shared Secure Agent groups

If you are the administrator of a parent organization, you can share a Secure Agent group with the sub-organizations. When you share a Secure Agent group, all sub-organizations can run data integration jobs on the Secure Agents within the group.

Note: Share a Secure Agent group when all agents in the group run only the Data Integration Server service. You cannot run non-data integration jobs on a shared Secure Agent group.

Share a Secure Agent group to optimize the use of available Secure Agent resources. For example, your organization contains separate sub-organizations for departments in different time zones. Each sub-organization runs data integration tasks at different times of the day. If you create one Secure Agent group for each sub-organization, some Secure Agent groups might be used heavily at certain times of the day while others remain idle. To distribute the tasks more evenly, add the Secure Agents to a Secure Agent group, and share the Secure Agent group with the sub-organizations.

To share a Secure Agent group, you must have the appropriate license.

When you share a Secure Agent group, the group appears on the **Runtime Environments** page in all sub-organizations. The sub-organization administrators cannot view the Secure Agents within the group. They cannot perform management tasks on the group such as adding or deleting Secure Agents, renaming, deleting, or unsharing the group, or changing the group permissions.

When a user in the sub-organization creates a connection or task, the user can select the shared Secure Agent group as the runtime environment.

Flat file connections in shared Secure Agent groups

If a shared Secure Agent group contains multiple Secure Agents and the group is used as the runtime environment for a flat file connection, the directory used in the connection must be accessible by all Secure Agents in the group.

If the directory is not accessible by all Secure Agents, tasks that use the connection fail if they are assigned to a Secure Agent that cannot access the directory.

Working with Secure Agent groups

Create Secure Agent groups on the **Runtime Environments** page. After you create a Secure Agent group, you can rename or delete the group, add and remove Secure Agents, and change group permissions.

You can complete the following tasks:

Create a Secure Agent group.

To create a Secure Agent group, click **New Runtime Environment** and enter a name for the group. After you create a group, you can add Secure Agents to the group.

Rename a Secure Agent group.

To rename a Secure Agent group, expand the Actions menu, select **Rename Secure Agent Group**, and enter a new name for the group. Informatica Intelligent Cloud Services updates the group name in all services that use the group.

Enable or disable services for a Secure Agent group.

To enable or disable services for a Secure Agent group, expand the Actions menu, select **Enable or Disable Services**, and select the services to enable or disable. You can enable or disable any service that your organization is licensed to use.

Note: Before you disable a service, verify that no connection, task, or process that uses the group as the runtime environment requires the service. If a connection, task, or process has a Secure Agent group selected as the runtime environment and you disable a required service, the task or process cannot run. Similarly, if a feature has a Secure Agent group selected as the runtime environment and you disable a required service, the feature cannot be used.

Add Secure Agents to a group.

To add Secure Agents to a group, expand the Actions menu and select **Add or Remove Secure Agents**. You can add any agent that is in the Unassigned Agents group on the **Runtime Environments** page.

Alternatively, you can add a new Secure Agent to an existing group by setting the InfaAgent.GroupName property in the infaagent.ini file before you register the agent.

When you add more than one Secure Agent to a Secure Agent group, all agents must meet the following requirements:

- All of the agents must be of the same type, for example, all local agents or all agents that run on Amazon EC2 machines.
- Each Secure Agent must be configured to connect to the same external systems and have access to files such as libraries, initialization files, and JAR files.
- Each Secure Agent must have access to the files used in tasks. Ensure that all files used in a task are available in a shared location.

Remove Secure Agents from a group.

To remove Secure Agents from a group, expand the Actions menu and select **Add or Remove Secure Agents**. When you remove an agent from a group, Informatica Intelligent Cloud Services adds it to a group named "Unassigned Agents."

You can remove an agent from a Secure Agent group if the group is not used as the runtime environment for a connection or task. If the group is used, you can remove an agent if it is not the only agent in the group.

Delete a Secure Agent group.

To delete Secure Agent group, expand the Actions menu and select **Delete Secure Agent Group**. You can delete a Secure Agent group if it does not contain any Secure Agents.

If the Secure Agent group is associated with an elastic configuration and the elastic cluster is running, you must stop the cluster and associate the configuration with a different runtime environment before you can delete the group.

Share or unshare a Secure Agent group.

If you are the administrator of a parent organization, you can share a Secure Agent group so that the sub-organizations can use it. You can unshare a group if it is not used in a connection or task. From the Actions menu associated with the group, choose **Share Secure Agent Group** or **Unshare Secure Agent Group**.

Change permissions for a Secure Agent group.

To change permissions for a Secure Agent group, expand the Actions menu and select **Change Permissions**. You can define permissions for a Secure Agent group for each user group in your organization.

You can set the following permissions:

Permission	Description
Read	View details about the Secure Agent group and use the Secure Agent group in a task.
Update	Edit the Secure Agent group.
Delete	Delete the Secure Agent group.
Change	Change permissions for the Secure Agent group.

Adding a Secure Agent to a group

You can add any available Secure Agent to a Secure Agent group. Available agents appear in the "Unassigned Agents" group on the **Runtime Environments** page. You cannot add a Secure Agent to a group if the agent has already been added to another group.

1. In Administrator, select **Runtime Environments**.
2. Expand the Actions menu for the Secure Agent group, and select **Add or Remove Secure Agents**.
3. In the **Available Agents** list, enable the checkbox for the Secure Agents that you want to add to the group.

   If no agent names are enabled in the **Available Agents** list, then all agents are added to other groups. You must remove an agent from a group before you can add it to a different group.

   When you enable a checkbox, the Secure Agent appears in the **Selected Agents** list, as shown in the following image:

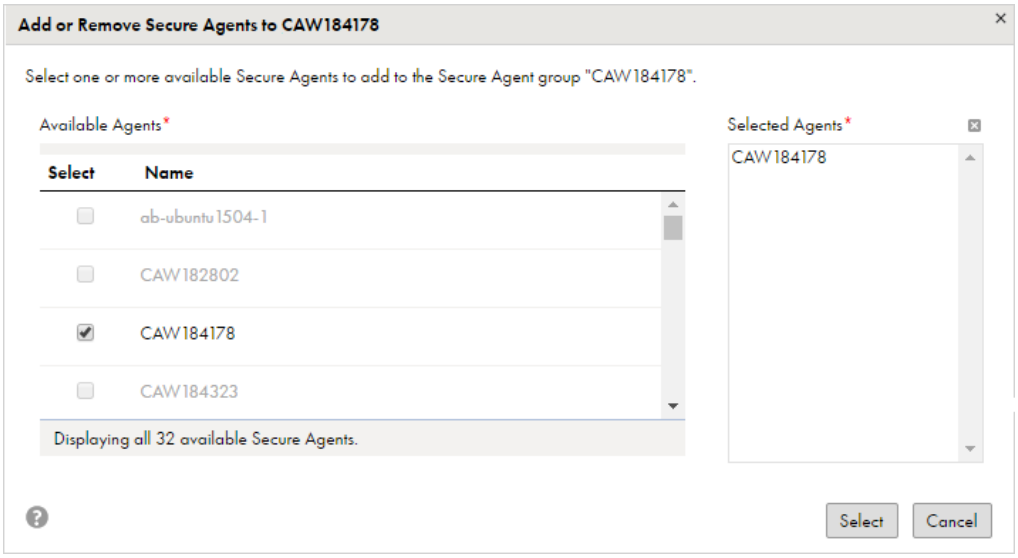

4. Click **Select.**

Adding a new Secure Agent to an existing group

You can add a Secure Agent to an existing Secure Agent group when you install the agent. To add a Secure Agent to an existing group, add the InfaAgent.GroupName property to the infaagent.ini file before you register the agent.

1. Install the Secure Agent.
2. On Windows, when you are prompted to register the agent, open Windows Services and stop the agent.

   On Linux, when the installation program finishes, do not start the agent.
3. Open `<Secure Agent installation directory>/apps/agentcore/conf/infaagent.ini` in a text editor.
4. Add the following property and save the file:

   ```
   InfaAgent.GroupName=<Secure Agent group name>
   ```
5. Start the agent.
6. Register the agent.

   Informatica Intelligent Cloud Services adds the Secure Agent to the group you specify in the InfaAgent.GroupName property instead of a new group.

Removing a Secure Agent from a group

You can remove an agent from a Secure Agent group if the group is not used in a connection or task. If the group is used in a connection or task, you can remove an agent if it is not the only agent in the group. When you remove a Secure Agent from a group, Informatica Intelligent Cloud Services adds it to a group named "Unassigned Agents."

1. In Administrator, select **Runtime Environments**.
2. Expand the Actions menu for the Secure Agent group, and select **Add or Remove Secure Agents**.
3. In the **Selected Agents** list, select the agents that you want to remove from the group, and click **X**.

   The check box for each agent that you remove is disabled and the Secure Agents no longer appear in the **Selected Agents** list, as shown in the following image:

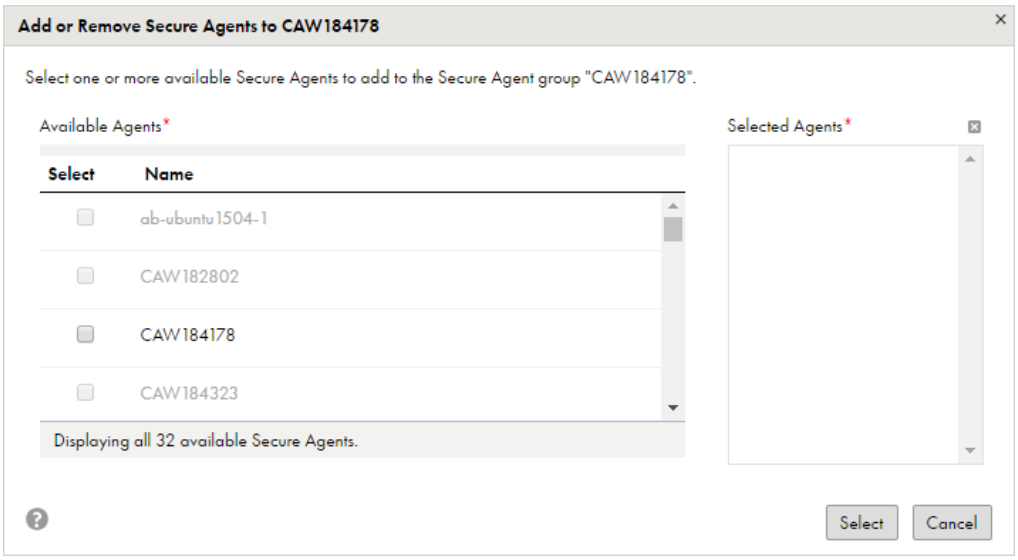

4. Click **Select.**

The Secure Agent appears in the "Unassigned Agents" group on the **Runtime Environments** page.

Viewing Secure Agent group dependencies

You can view object dependencies for Secure Agent groups.

When you view dependencies for a Secure Agent group, Administrator lists the connections and assets in each service that use the group as the runtime environment.

To view object dependencies for a Secure Agent Group, expand the Actions menu and select **Show Dependencies**.

The following image shows the **Dependencies** page for a Secure Agent group:

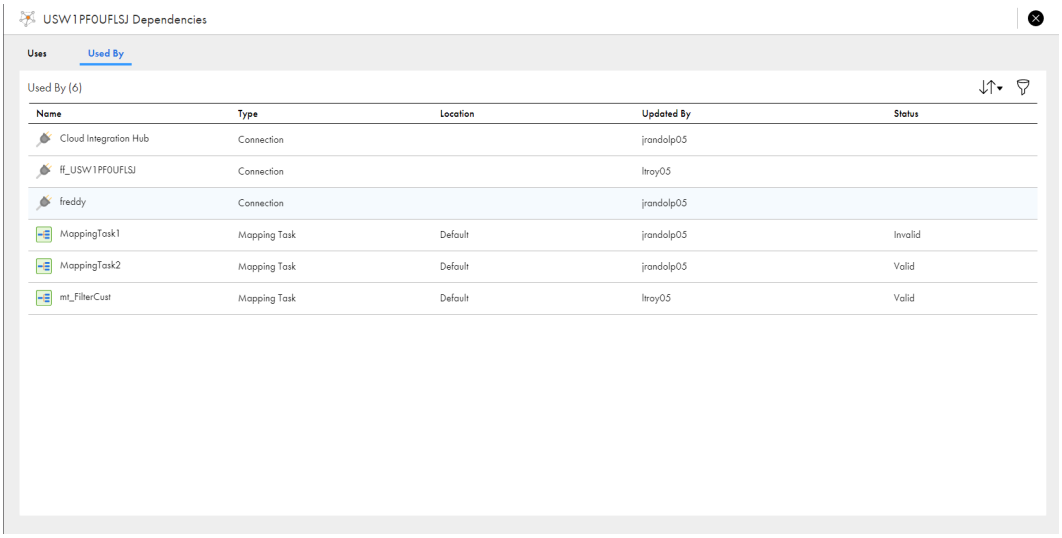

To sort the objects that appear on the page, click the sort icon and select the column name for the property you want to sort by.

To filter the objects that appear on the dependencies page, click the Filter icon. Use filters to find specific objects. To apply a filter, click **Add Field**, select the property to filter by, and then enter the property value. You can specify multiple filters. For example, to find connections with Oracle in the name, add the Type filter and specify Connection. Then add the Name filter and enter "Oracle."

Secure Agents

The Informatica Cloud Secure Agent is a lightweight program that runs all tasks and enables secure communication across the firewall between your organization and Informatica Intelligent Cloud Services. When the Secure Agent runs a task, it connects to the Informatica Cloud hosting facility to access task information. It connects directly and securely to sources and targets, transfers data between them, orchestrates the flow of tasks, runs processes, and performs any additional task requirement.

If the Secure Agent loses connectivity to Informatica Intelligent Cloud Services, it tries to reestablish connectivity to continue the task. If it cannot reestablish connectivity, the task fails.

The Secure Agent uses pluggable microservices for data processing. For example, the Data Integration Server runs all data integration jobs, and Process Server runs application integration and process orchestration jobs. Each service has a unique set of configuration properties, such as Tomcat and Tomcat JRE settings. For more information about Secure Agent services, see Chapter 13, “Secure Agent services” on page 114.

You can install and run one Secure Agent on a physical or virtual machine. After you install a Secure Agent, all users in the organization share the Secure Agent. You can configure the Secure Agent properties and move it to a different Secure Agent group. To improve scalability, you can also add multiple agents to a Secure Agent group.

Working with Secure Agents

After you create a Secure Agent, you might need to perform management tasks such as viewing and configuring agent properties, checking the host information, viewing audit logs, or refreshing the agent status. You can also delete a Secure Agent if it is no longer used.

You perform most management tasks for Secure Agents on the agent details page. To access the agent details page, click a Secure Agent on the **Runtime Environments** page.

The following image shows the agent details page:

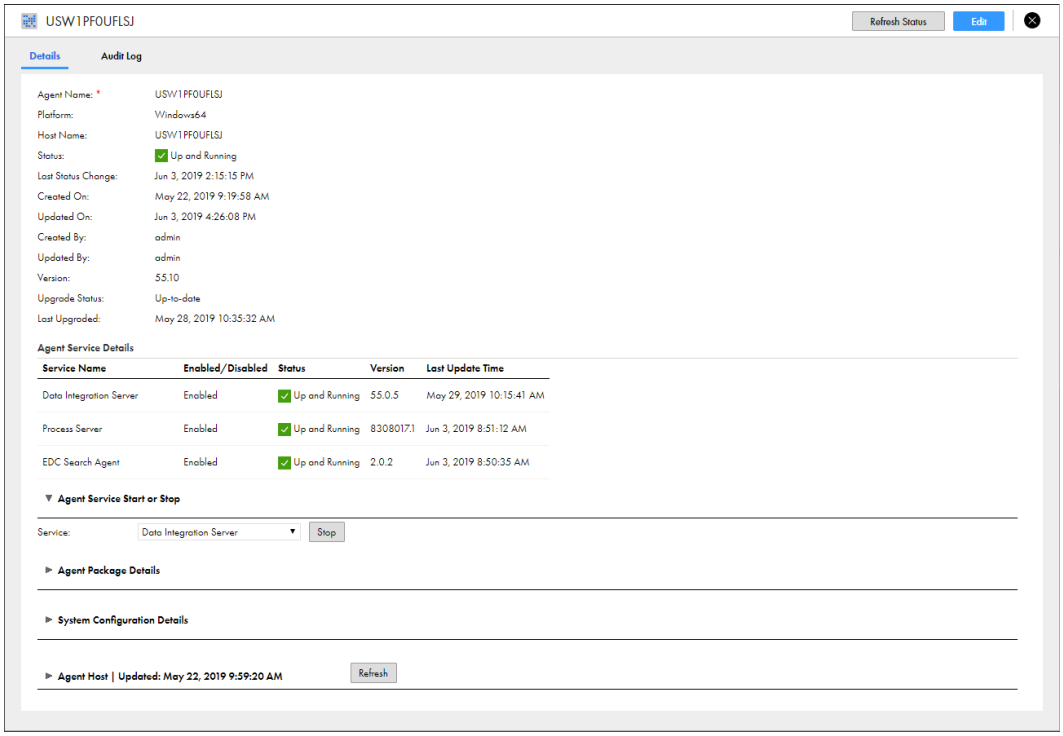

You can complete the following tasks:

View the Secure Agent details.

View details such as the host name, the current status, the last date and time that the agent was updated, and the agent version.

The Secure Agent can have any of the following statuses:

Status	Description
Agent Core is not running.	The Secure Agent is not available, but one or more of the services is running.
Not all the services are running.	The Secure Agent is available, but one or more of the services is not available.
Agent Core Upgrading	The Secure Agent is upgrading to a new version.
Stopped	The Secure Agent is not available.
Up and Running	The Secure Agent and all of the services that the agent runs are available.

View the agent service details.

View details for services that run on the Secure Agent such as the service name, status, version, and last update time.

A service can have any of the following statuses:

Status	Description
Error	The process failed.
Restarting Due to Error	The service is starting due to a failure.
Shutting Down	The service is shutting down.
Standby	The service is running, but it is not compatible with Informatica Intelligent Cloud Services.
Starting Up	The service is starting up.
Stopped	The service is not available.
Up and Running	The service is running.
Warning	The service is running, but it cannot accept work.

The version number changes each time you modify the service. The Secure Agent retains the directories for the old version of the service for seven days. For example, if you update the NetworkTimeoutPeriod for version 55.0.2 of the Data Integration Server, the agent increments the version number to 55.0.3 and creates the following directory:

```
<Secure Agent installation directory>/apps/Data_Integration_Server/55.0.3.1
```

It deletes the `<Secure Agent installation directory>/apps/Data_Integration_Server/55.0.2.x` directories after seven days.

Stop and start the services that run on the Secure Agent.

Stop and start the services that run on a Secure Agent to perform troubleshooting, optimize resources on the agent machine, or make service configuration changes. When you stop or start a service, other services that run on the agent are not affected.

View the Secure Agent package details.

Expand the **Agent Package Details** section to see the name and version number for the packages in each service that runs on the Secure Agent. You can filter the packages by service.

View and edit Secure Agent service properties.

Expand the **System Configuration Details** section to see the Secure Agent service properties. You can filter the properties by service and type.

To configure the properties, click **Edit**. You can configure properties for each service that runs on the Secure Agent. You can also add and remove custom properties, which are used by connectors. For more information about Secure Agent services and service properties, see Chapter 13, “Secure Agent services” on page 114. For more information about custom properties, see the help for the appropriate connector.

View the Secure Agent host properties.

Expand the **Agent Host** section to see information about the machine that hosts the Secure Agent. For example, you can view the machine name, operating system, and available disk space.

To refresh the information, click **Refresh**. The last date and time that the information was refreshed appears next to the **Agent Host | Updated** heading.

View the Audit Log.

To view audit information such as start and stop times, server connections, and upgrade messages, click **Audit Log**.

Refresh the Secure Agent status.

To refresh the status of the Secure Agent, click **Refresh Status** in the upper right corner of the page.

To view the status on Linux, you can also navigate to the following directory:

```
<Secure Agent installation directory>/apps/agentcore
```

Then run one of the following commands:

```
./consoleAgentManager.sh getstatus
```

```
./consoleAgentManager.sh updatestatus
```

Stopping and starting services on a Secure Agent

By default, each Secure Agent in an organization runs all microservices that are used for data processing in the organization. You can stop and start these services to perform troubleshooting, optimize resources on the agent machine, or make configuration changes. When you stop or start a Secure Agent service, other services that run on the agent are not affected.

The services that you stop and start on a Secure Agent are the Secure Agent services, which are different from the Informatica Intelligent Cloud Services. For example, if you want to stop the services associated with Operational Insights, you must stop the OI Data Collector service on the agent. For more information about Secure Agent services, see Chapter 13, “Secure Agent services” on page 114.

You might need to stop and restart a Secure Agent service in the following circumstances:

You need to troubleshoot issues with a specific service.

If a service shows an error state, you can stop the service, troubleshoot the problem, and then restart the service.

You are running memory or CPU intensive jobs, and you want to optimize computing resources on the Secure Agent machine.

For example, your organization runs Data Integration and Application Integration jobs. You want to optimize computing resources so that the Data Integration jobs run during the day and the Application Integration jobs run at night. To do this, stop Process Server during the day and restart it in the evening, and stop the Data Integration Server at night and restart it in the morning.

You update service configuration properties for the File Integration Service.

After you change configuration properties for the File Integration Service, you must restart the service. If the Secure Agent runs other services, you can stop and restart the File Integration Service without affecting the other services.

To start or stop a service on a Secure Agent, you must have update permission on the Secure Agent.

If you are the administrator of a sub-organization, you can start and stop services on the agents in the sub-organization. However, you cannot start and stop services on a Secure Agent that is in a shared Secure Agent group.

Each time you start and restart a service, the Secure Agent creates a new subdirectory for the service-related files. For example, if the Secure Agent uses version 12.1 of the B2B Processor Service, the Secure Agent installation directory contains the following subdirectory:

```
<Secure Agent installation directory>/apps/B2BProcessor/12.1.1
```

When you stop and restart the B2B Processor service, the Secure Agent creates the following directory:

```
<Secure Agent installation directory>/apps/B2BProcessor/12.1.2
```

The Secure Agent does not delete the `.../12.1.1` directory.

Example

Your organization uses Data Integration and has licenses for Enterprise Data Catalog integration, file integration, and mass ingestion.

Your Secure Agent runs the following services:

- Data Integration Server
- EDC Search Agent
- File Integration Service
- Mass Ingestion

If you have issues with Enterprise Data Catalog search, you can stop the EDC Search Agent service while you perform troubleshooting. When you stop the EDC Search Agent service, you cannot perform data catalog discovery in Data Integration. However, jobs processed by the other services on this agent such as mappings, tasks, taskflows, and AS2 file transfers continue to run.

Guidelines for stopping and starting services

Use the following guidelines when you stop and start services on a Secure Agent:

- Use caution when you stop services on a Secure Agent because stopping services can cause job failures.

  When you stop a service, any job that requires the service and is currently running on the agent stops. If there are no other agents in the group, the job can no longer run. If there are other agents in the group, you can restart the job and it will run on a different agent.
- Do not stop the Data Integration Server on an agent if you store connection properties with the agent.

  If you store connection properties with a local Secure Agent and you stop the Data Integration Server on the agent, users will not be able to access any connections or run tasks in the organization. Any job that is currently running on the agent also fails.
- Do not stop and start services to reserve a Secure Agent group for certain types of jobs.

  If you want to reserve a Secure Agent group for certain types of jobs, you can enable the required services for the Secure Agent group and disable other services. For more information about enabling and disabling services for a Secure Agent group, see "Service assignment for Secure Agent groups" on page 86.

Stopping a Service

You can stop a service that is in the "Up and Running" or "Error" state. Stopping a service stops all versions of the service that are running. After a service stops, you can start the latest version of the service.

Note: If you stop a service and then restart the Secure Agent, the service remains stopped until you start it.

1. In Administrator, select **Runtime Environments**.
2. On the **Runtime Environments** page, click the name of the Secure Agent.

   Note: You might have to expand the Secure Agent group to see the list of Secure Agents within the group.
3. Click the **Details** tab.
4. In the **Agent Service Start or Stop** area, select the service that you want to stop.
5. Click **Stop**.

The service stops, and Informatica Intelligent Cloud Services adds an entry in the audit log indicating that the service was stopped by a user.

Starting a Service

You can start a service that is in the "Stopped" state. Starting a service starts the latest version of the service.

1. In Administrator, select **Runtime Environments**.
2. On the **Runtime Environments** page, click the name of the Secure Agent.

   Note: You might have to expand the Secure Agent group to see the list of Secure Agents within the group.
3. Click the **Details** tab.
4. In the **Agent Service Start or Stop** area, select the service that you want to start.
5. Click **Start**.

Informatica Intelligent Cloud Services attempts to start the service. After the service starts, the status changes to "Up and Running." If the service fails to start, check the audit log to find the cause of the error.

Configuring agent blackout periods

You can configure blackout periods for a Secure Agent. Blackout periods prevent data integration jobs from running on the agent during a certain period. Configure an agent blackout period to configure specific hours, days, or intervals in which no data integration jobs can run on the agent.

Agent blackout periods stop the Data Integration Server service from running jobs on a Secure Agent during the blackout period. They do not prevent other types of jobs from running on the agent. Configure an agent blackout period in the following circumstances:

- The Data Integration Server is the only service enabled on the agent and you want to stop all data integration jobs from running during a certain period.
- The Secure Agent runs multiple services and you want to stop only the data integration jobs from running during a certain period.

Note: The agent blackout period is different than the schedule blackout period for the organization. During an organization's schedule blackout period, no jobs can run on any agent. For more information about schedule blackout periods, see “Configuring a blackout period” on page 154.

To configure a blackout period on a Secure Agent, you must create a blackout file. The blackout file is an XML file that specifies the repeat frequency, start date, and end date for each blackout period.

For example, the following blackout file contains two blackout periods: one blackout period from July 27, 5:00 AM through July 28, 11:00 PM and a second blackout period that repeats on Fridays from 2:00-4:00 PM:

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<BlackoutWindows>
    <BlackoutWindow>
        <RepeatFrequency>OneTime</RepeatFrequency>
        <Start>2019-07-27 5:00:00</Start>
        <End>2019-07-28 23:00:00</End>
    </BlackoutWindow>
    <BlackoutWindow>
        <RepeatFrequency>Friday</RepeatFrequency>
        <Start>14:00:00</Start>
        <End>16:00:00</End>
    </BlackoutWindow>
</BlackoutWindows>
```

To configure one or more blackout periods, create a file named "blackoutWindows.dat" in the following directory:

```
<Secure Agent Installation Directory>\apps\Data_Integration_Server\conf\
```

If you want to use a different file name and directory, you can override the file name and file path.

After you create a blackout file, restart the Data Integration Server service on the Secure Agent so that the blackout periods take effect.

Overriding the blackout file name and directory

You can override the blackout file name and directory.

To do this, set the following custom property for the Data Integration Server on the agent details page:

Service	Type	Name	Value
Data Integration Server	Tomcat	BlackoutWindowsFile	File path and file name for the blackout file. For example: `C:/AgentBlackouts/Agent001Blackouts.dat` **Note:** Use forward slashes (/) in the file path on both Windows and UNIX machines because the Secure Agent interprets backslashes (\) as escape characters. The file path must be accessible by the Secure Agent.

For more information about configuring custom properties for a Secure Agent service, see "Configuring Secure Agent service properties" on page 145.

Blackout file structure

The blackout file is an XML file that contains elements that define each blackout period and the frequency, start time, and end time for each blackout period.

The blackout file has the following structure:

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<BlackoutWindows>
    <BlackoutWindow>
        <RepeatFrequency></RepeatFrequency>
        <Start></Start>
        <End></End>
    </BlackoutWindow>
    <BlackoutWindow>
        <RepeatFrequency></RepeatFrequency>
        <Start></Start>
        <End></End>
```

```
        </BlackoutWindow>
        …
    </BlackoutWindows>
```

The file contains the following elements:

Element	Required/ Optional	Description
BlackoutWindows	Required	Contains a BlackoutWindow element for each blackout period. Must contain one or more BlackoutWindow elements.
BlackoutWindow	Required	Defines one blackout period. Must contain one RepeatFrequency element, one Start element, and one End element.
RepeatFrequency	Required	Repeat frequency for the blackout period. Must contain one of the following values: - OneTime - Daily - Weekdays - Sunday - Monday - Tuesday - Wednesday - Thursday - Friday - Saturday
Start	Required	Blackout period start time in the format yyyy-mm-dd hh24:mi:ss. For example, 2019-07-25 10:26:55. The time zone is the Secure Agent time zone.
End	Required	Blackout period end time in the format yyyy-mm-dd hh24:mi:ss. For example, 2019-07-26 11:45:00. The time zone is the Secure Agent time zone.

Do not enclose element values in quotation marks.

Renaming a Secure Agent

By default, the name of a Secure Agent is the same as the name of the machine where you installed the agent. You can change the agent name.

1. On the **Runtime Environments** page, click the name of the Secure Agent.

   Note: You might have to expand the Secure Agent group to see the list of Secure Agents within the group.
2. Click the **Details** tab.
3. In the upper right corner, click **Edit**.
4. Enter a new name in the **Agent Name** field.
5. Click **Save**.

Deleting a Secure Agent

Delete a Secure Agent if you no longer need it to run tasks. Delete a Secure Agent on the **Runtime Environments** page.

Note: You cannot delete a Secure Agent if it is used in a connection or a task. For example, if the Secure Agent is the only agent in a group, and the group is used as the runtime environment for a connection or task, you cannot delete the agent.

1. In Administrator, select **Runtime Environments**.
2. Expand the Actions menu for the Secure Agent and select **Delete Secure Agent**.

   If the Secure Agent is running, a warning message appears. Stopping an active Secure Agent prevents scheduled tasks associated with the Secure Agent from running. Ignore the warning if you do not need the Secure Agent.

If you no longer need the Secure Agent, uninstall the Secure Agent after you delete it.

Upgrading a Secure Agent

The Secure Agent upgrades automatically the first time that you access a new Informatica Intelligent Cloud Services release. The upgrade process installs a new version of the Secure Agent, updates connector packages, and applies configuration changes for the microservices that run on the agent. You do not need to upgrade the Secure Agent manually.

However, to prepare for an upgrade, you might need to perform tasks such as ensuring that each Secure Agent machine has enough disk space available for the upgrade. For more information about preparing for an upgrade, see *Administrator What's New*.

Secure Agent Manager

When you install the Secure Agent on Windows, you also install the Informatica Cloud Secure Agent Manager. The Secure Agent runs as a Windows service. You can launch the Secure Agent Manager from the Windows Start menu or the desktop icon.

Use the Secure Agent Manager to perform the following tasks:

- View the status of the Secure Agent and the services that the Secure Agent runs.
- Stop and restart the Secure Agent.
- Configure Windows settings such as proxy settings and a Windows Secure Agent service login.

The Secure Agent Manager displays the status of the Secure Agent and the services that the Secure Agent runs. If the Secure Agent or one of the services that the Secure Agent runs is not starting or not running, the Secure Agent Manager displays an alert message and a link that you can click to view details.

When you close the Secure Agent Manager, it minimizes to the Windows taskbar for quick access. Closing the Secure Agent Manager does not stop the Secure Agent. When the Secure Agent Manager is minimized, you can view the Secure Agent status by hovering over the Secure Agent Manager icon.

Configuring a proxy to exclude non-proxy hosts

A proxy server allows indirect connection to network services for security and performance reasons. For example, you can use a proxy server to get through a firewall, and some proxies provide caching

mechanisms. When you configure a proxy server for the Informatica Cloud Secure Agent, you can exclude certain IP addresses and host names from the proxy.

When you configure a proxy server for the Informatica Cloud Secure Agent, you define the minimum required settings in the Secure Agent Manager. Informatica Intelligent Cloud Services updates the following file and adds other properties that you can edit manually:

```
<Secure Agent installation directory>/apps/agentcore/conf/proxy.ini
```

The property, InfaAgent.NonProxyHost, enables you to exclude IP addresses or host names. By default, Informatica Intelligent Cloud Services adds localhost as the value for InfaAgent.NonProxyHost when you initially configure the proxy server:

```
InfaAgent.ProxyPassword=ZU8KjIzgtVrVmfRMUPzPMw\=\=
InfaAgent.ProxyNtDomain=
InfaAgent.ProxyHost=foo.bar.com
InfaAgent.ProxyPasswordEncrypted=true
InfaAgent.NonProxyHost=localhost|127.*|[\:\:1]
InfaAgent.ProxyUser=
InfaAgent.ProxyPort=12345
InfaAgent.AuthenticationOrder=
```

To exclude certain IP addresses or host names from the proxy, perform the following steps:

1. Open `<Secure Agent installation directory>/apps/agentcore/conf/proxy.ini`.
2. Update the value for InfaAgent.NonProxyHost to specify the IP addresses or host names that you want to exclude.

   For example:

   - Local IP addresses:

     ```
     InfaAgent.NonProxyHost=localhost|127.|[\:\:1]|123.432.
     ```

   - Host names:

     ```
     InfaAgent.NonProxyHost=localhost|127.|[\:\:1]|.foo.com
     ```

   Note: You can combine a list of host names and IP addresses using the pipe character (|) as a delimiter. You can also enter a wildcard to the left for host names or to the right for IP addresses.
3. Restart the Secure Agent so that the changes take effect.

Stopping and restarting the Secure Agent on Windows

The Secure Agent Manager displays the Secure Agent status. You can use the Secure Agent Manager to stop or restart the Secure Agent.

Launch the Secure Agent Manager from the Windows **Start** menu. If the Secure Agent Manager is active, you can click the Informatica Cloud Secure Agent Manager icon in the Windows taskbar notification area to open the Secure Agent Manager.

To stop the Secure Agent from the Secure Agent Manager, click **Stop**. To restart the Secure Agent, click **Restart**. The Secure Agent Manager displays a message when the action is complete.

When you close the Secure Agent Manager, it minimizes to the Windows taskbar notification tray. Closing the Secure Agent Manager does not stop the Secure Agent.

Starting and stopping the Secure Agent on Linux

After you download the Secure Agent program files on a Linux machine, you can run the Secure Agent as a Linux process. Manually start the Secure Agent process on Linux.

1. From the command line, navigate to the following directory:

   ```
   <Secure Agent installation directory>/apps/agentcore
   ```

2. To start the Secure Agent, enter the following command:

   ```
   ./infaagent startup
   ```

3. To stop the Secure Agent, enter the following command:

   ```
   ./infaagent shutdown
   ```

You can view the Secure Agent status from Informatica Intelligent Cloud Services or from a Linux command line.

CHAPTER 12

Serverless runtime environments

A serverless runtime environment is an advanced serverless deployment solution that doesn't require downloading, installing, configuring, and maintaining a Secure Agent or Secure Agent group. You can use a serverless runtime environment in the same way that you use a runtime environment when you configure a connection or some types of tasks in Data Integration.

Compared to the multi-tenant model on the Hosted Agent, a serverless runtime environment uses an isolated, single-tenant model. The model provides a dedicated server with virtual machine resources to run tasks for your organization. The serverless runtime environment auto-scales with the size of the workload while your data remains in your cloud environment.

A serverless runtime environment is hosted in Informatica's Amazon Virtual Private Cloud (VPC) on the AWS cloud platform. The serverless runtime environment creates a cross-account elastic network interface (ENI) to connect to your cloud environment.

Note: To use a serverless runtime environment, your cloud environment must be on the AWS cloud platform and your VPC must have default tenancy. A serverless runtime environment cannot connect to a VPC with dedicated instance tenancy.

To use a serverless runtime environment, your organization must have the appropriate licenses.

Using a serverless runtime environment for Data Integration Elastic

When you use a serverless runtime environment to run elastic jobs, the advanced serverless deployment of Data Integration Elastic is configured with the prerequisites to create an elastic cluster and to run the jobs on the cluster.

The serverless runtime environment manages the elastic cluster while the cluster adapts to workload changes by provisioning and deprovisioning resources.

Serverless compute units

Serverless compute units represent CPUs and memory that a serverless runtime environment can use to run tasks.

When you create a serverless runtime environment, you configure a maximum number of serverless compute units that each task can request from the serverless runtime environment. When you create a mapping task, you can override the maximum number of compute units that the task can request. In Monitor, you can view the number of compute units that the task requested and consumed.

If the task runs longer than the task timeout that you specify, the serverless runtime environment terminates the task.

For information about the meter, see “Metering serverless compute units” on page 40.

Before you begin

Before you can create a serverless runtime environment, you must set up your cloud environment to connect to the serverless runtime environment.

Complete the following tasks:

1. Optionally, create a NAT gateway that the serverless runtime environment can use to connect to external services.
2. Optionally, create S3 folders for supplementary files, such as JAR files and external libraries.
3. Set up an IAM role that can be used to create an ENI.
4. Create a security group that the serverless runtime environment will attach to the ENI.

Step 1. Create a NAT gateway

Optionally, create a NAT gateway to allow the subnet that is configured in the serverless environment to connect to external services through the internet.

You must create a NAT gateway in the following situations:

- Tasks access Amazon S3 sources and targets that are in a different AWS region.
- Tasks access sources and targets that are not on AWS.

When you configure the NAT gateway, complete the following tasks, configure the NACL (network access control list) that is associated with the subnet with inbound rules to allow all traffic on the following ports:

- Ephemeral port range `1024-65535`
- Port 443

For information about creating a NAT gateway, refer to the AWS documentation.

Step 2. Create S3 folders for supplementary files

If your environment and data integration jobs require JAR files and external libraries, dedicate a location on Amazon S3 to store the files and create folders for each file type. The serverless runtime environment will access the location to retrieve the files.

Create the following file structure on Amazon S3:

```
S3 location for supplementary files
├── ext
│   └── python
├── odbc
│   └── lib
└── jars
    └── ctjars
```

Store the following types of files in each location:

Location	Files
`<S3 location>/ext`	JDBC JAR files
`<S3 location>/ext/python/`	Python installation and resource files used in the Python transformation

Location	Files
<S3 location>/odbc	The following files: - odbc.ini - odbcinst.ini - exports.ini
<S3 location>/odbc/lib	ODBC shared libraries for a Linux operating system
<S3 location>/jars/ctjars	JAR files to use with the Java transformation

Step 3. Set up an IAM role

Create an IAM role to establish trust between your AWS account and the Informatica AWS account so that the serverless runtime environment can create an ENI and securely connect to data sources in your cloud environment.

Create a cross-account IAM role in your AWS account that identifies Informatica as a trusted entity.

1. Create a role for another AWS account.
2. In the trust relationship, specify the Informatica account number and the external ID.

   For example, specify the following policy in the trust relationship:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws:iam::<Informatica account>:root"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "sts:ExternalId": "<External ID>"
        }
      }
    }
  ]
}
```

3. Edit the role permissions and specify a policy to grant the serverless runtime environment a minimal set of permissions on your account.

   Use the following template for the policy:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "VisualEditor0",
            "Effect": "Allow",
            "Action": [
                "ec2:DetachNetworkInterface",
                "ec2:DeleteTags",
                "ec2:DescribeTags",
                "ec2:CreateTags",
                "ec2:DeleteNetworkInterface",
                "ec2:DescribeSecurityGroups",
                "ec2:CreateNetworkInterface",
                "ec2:DeleteNetworkInterfacePermission",
                "ec2:DescribeNetworkInterfaces",
```

```
                "ec2:DescribeAvailabilityZones",
                "ec2:CreateNetworkInterfacePermission",
                "ec2:AttachNetworkInterface",
                "ec2:DescribeNetworkInterfacePermissions",
                "ec2:DescribeSubnets",
                "ec2:DescribeNetworkAcls"
            ],
            "Resource": "*"
        },
        {
            "Sid": "VisualEditor1",
            "Effect": "Allow",
            "Action": [
                "s3:PutObject",
                "s3:GetObject",
                "s3:ListBucket",
                "s3:DeleteObject",
                "s3:GetBucketAcl"
            ],
            "Resource": [
                "arn:aws:s3:::<S3 location for supplementary files>",
                "arn:aws:s3:::<S3 location for supplementary files>/*"
            ]
        }
    ]
}
```

For more information about setting up cross-account IAM roles, refer to the AWS documentation.

Step 4. Create a security group

Create a security group that the serverless runtime environment will attach to the ENI.

When you create the security group, verify that the security group restricts all inbound traffic and permits all outbound traffic.

Serverless runtime environment properties

Create a new serverless runtime environment and configure properties on the **Serverless Environments** page. To view properties for a serverless runtime environment, expand the **Actions** menu for the environment and select **View**.

The following image shows the **Serverless Environments** page:

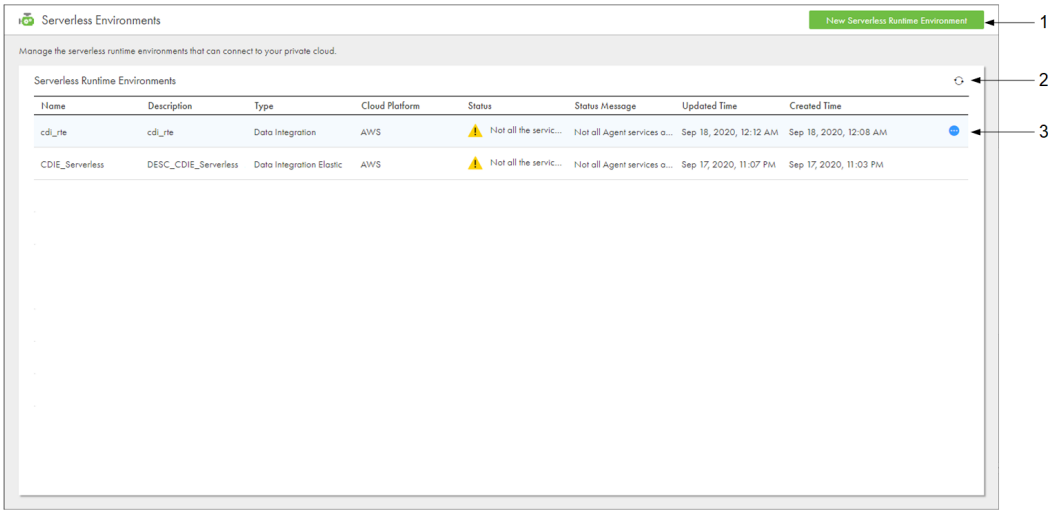

1. Option to create a new serverless runtime environment
2. Refresh icon
3. Actions menu

Basic configuration

The following table describes the basic properties that you configure for the serverless runtime environment:

Property	Description
Name	Name of the serverless runtime environment.
Description	Description of the serverless runtime environment.
Task Type	Type of tasks that run in the serverless runtime environment. - Select Data Integration to run tasks that are created using Data Integration. - Select Data Integration Elastic to run tasks that are created using Data Integration Elastic.
Cloud Platform	Cloud platform to host the serverless runtime environment. You can use only Amazon Web Services (AWS).
Max Compute Units	Maximum number of serverless compute units corresponding to machine resources that a task can use.
Task Timeout	Amount of time in minutes to wait for a task to complete before it is terminated. The timeout ensures that serverless compute units are not unproductive when a task hangs. By default, the timeout is 2880 minutes (48 hours). You can set the timeout to a value that is less than 2880 minutes.

Cloud Data Integration configuration

The following table describes the properties that you use to create an IAM role for the serverless runtime environment:

Property	Description
Informatica Account Number	Informatica's account number on the cloud platform where the serverless runtime environment will be created. The account number is populated automatically.
External ID	External ID to associate with the role that you create for the serverless runtime environment. You can use the generated external ID or specify your own external ID.

AWS resource configuration

The AWS resource configuration specifies how the serverless runtime environment connects to your AWS account and the sources and targets that you use in mappings.

The following table describes the properties:

Property	Description
Configuration Name	Name of the AWS resource configuration. The name must consist of ASCII characters up to a maximum of 255 characters.
Configuration Description	Description of the AWS resource configuration. The description must consist of ASCII characters up to a maximum of 255 characters.
Account Number	Your account number on the cloud platform.
Region	Region on your cloud platform. The sources and targets that you use in mappings must either reside or be accessible from the region.
AZ ID	Identifier for the availability zone. The sources and targets that you use in mappings must either reside or be accessible from the availability zone.
VPC ID	ID of the Amazon Virtual Private Cloud (VPC). The VPC must be configured with an endpoint to access the sources and targets that you use in mappings. For example, `vpc-2f09a348`.
Subnet ID	ID of the subnet within the VPC. The subnet must be have an entry point to access the sources and targets that you use in mappings. For example, `subnet-b46032ec`.
Security Group ID	ID of the security group that the serverless runtime environment will attach to the ENI. The security group allows access to the sources and targets that you use in tasks. For example, `sg-e1fb8c9a`.
Role Name	Name of the IAM role that the serverless runtime environment can assume on your AWS account. The role must have permissions to create, read, delete, list, detach, and attach an ENI. It also requires read and write permissions on supplementary file location. Use the Informatica account number and the external ID when you create a policy for the role.

Property	Description
AWS Tags	AWS tags to label the ENI that is created in your AWS account. Each tag must be a key-value pair in the format: `Key=string,Value=string` where `Key` and `Value` are case-sensitive. Use a space to separate tags. **Note:** Follow the rules and guidelines for tagging that AWS specifies. For more information, refer to the AWS documentation.
Supplementary File Location	Location on Amazon S3 to store supplementary files, such as JAR files and external libraries for certain transformations and connectors. For example, if you use third-party or custom Java packages in the Java transformation, add the JAR files to an S3 location and specify the S3 location for the serverless runtime environment. Use the format: `s3://<bucket name>/<folder name>`.

Editing a serverless runtime environment

You can edit properties in a serverless runtime environment based on the environment status:

- Up and Running. You can update the maximum number of serverless compute units or the task timeout.

  If you edit the maximum number of serverless compute units or the task timeout, the updated values take effect for subsequent task runs.
- Failed. You can update any of the properties. Redeploy the serverless runtime environment for the properties to take effect.

If the serverless runtime environment has a different status, you must delete the serverless runtime environment and create a new one to edit the properties.

Before you delete a serverless runtime environment, complete the following tasks:

- Use Monitor to make sure that the environment is not running any jobs.
- Remove the serverless runtime environment from any connections and tasks that use it.

Redeploying a serverless runtime environment

You might redeploy the serverless runtime environment in the following situations:

- You change your licenses.
- The serverless runtime environment shut down because the organization ran out of serverless compute units. You can add more compute units to your organization and redeploy the serverless runtime environment.
- You update the configuration in your cloud environment. For example, you update the JAR files in the supplementary file location, or you update the policy that is attached to the IAM role.

Before you redeploy a serverless runtime environment, use Monitor to make sure that the environment is not running any jobs. Then, in Administrator, expand the **Actions** menu for the serverless runtime environment and click **Redeploy**.

Cloning a serverless runtime environment

You might clone a serverless runtime environment to create another environment that has a similar configuration. For example, you might want to create a similar serverless runtime environment that connects to a different subnet in your cloud environment or uses a different security group.

To clone a serverless runtime environment, you can expand the **Actions** menu for the serverless runtime environment and click **Clone**.

Rules and guidelines

Consider the following rules and guidelines when you create a serverless runtime environment:

- You can create a maximum of 10 serverless runtime environments in your organization. If you have a trial license, you can create a maximum of two environments.
- A serverless runtime environment can run a maximum of 10 tasks at the same time.
- It takes at least five minutes for the serverless runtime environment to become available. Use the **Serverless Environments** page to track the status of the environment and review any status messages.

Disaster recovery

If a disaster impacts the region or the availability zone that hosts a serverless runtime environment, redirect jobs to a temporary serverless runtime environment in a stable region or availability zone as part of your organization's disaster recovery plan.

Disaster recovery procedure

During a disaster, all virtual machines in the serverless runtime environment shut down and jobs can no longer run in the environment.

To minimize data loss and downtime, complete the following tasks:

1. Create a temporary serverless runtime environment in a stable region or availability zone.
2. Make sure that the connections used in jobs are available in the stable region or availability zone.
3. Clean up data related to incomplete job runs. If data was partially loaded to a target, manually delete the data or update the mapping to truncate the target before writing new rows.
4. Redirect jobs to the temporary environment.

Restoring the primary environment

When the region or availability zone that hosts the primary serverless runtime environment has recovered, you can restore the primary environment.

To restore the primary environment, complete the following tasks:

1. Clean up the ENIs that were created in your AWS account for the primary environment.
2. Redeploy the primary environment.
3. Redirect jobs to the primary environment.
4. Delete the temporary environment.

Connectors in a serverless runtime environment

The connectors that you can use with a serverless runtime environment depend on the type of mappings that the environment runs.

Use connectors based on one of the following mapping types:

Elastic mappings

The serverless runtime environment can connect to sources and targets using the following connectors:

- Amazon Redshift V2
- Amazon S3 V2
- JDBC V2
- Snowflake Cloud Data Warehouse V2

Mappings

The serverless runtime environment can connect to sources and targets using the following connectors:

- Amazon Aurora Connector
- Amazon Redshift V2 Connector
- Amazon S3 V2 Connector
- Box Connector
- Box Oauth Connector
- CDM Folders Connector
- Concur V2 Connector
- Coupa V2 Connector
- DB2 Warehouse on Cloud Connector
- Eloqua Bulk API Connector
- Google Analytics Connector
- Google BigQuery V2 Connector
- Google Cloud Spanner Connector
- Google Cloud Storage V2 Connector
- Marketo V3 Connector
- Microsoft Azure Blob Storage V3 Connector
- Microsoft Azure Cosmos DB SQL API Connector
- Microsoft Azure Data Lake Store Gen2 Connector
- Microsoft Azure Data Lake Store V3 Connector
- Microsoft Azure SQL Data Warehouse V3 Connector
- Microsoft Dynamics 365 for Sales Connector
- Microsoft Dynamics CRM Connector
- Microsoft SQL Server Connector
- MongoDB Connector
- MySQL Connector
- NetSuite V2 Connector

- Oracle Connector
- PostgreSQL Connector
- REST V2 Connector
- Salesforce Connector
- Salesforce Marketing Cloud Connector
- Salesforce Oauth Connector
- ServiceNow Connector
- Snowflake Cloud Data Warehouse V2 Connector
- SuccessFactors ODATA Connector
- SuccessFactors SOAP Connector
- Workday V2 Connector
- Zendesk V2 Connector

Note: The serverless runtime environment usage is specific to connectors. For more information, see the help for the relevant connector.

CHAPTER 13

Secure Agent services

Secure Agent services are pluggable microservices that the Secure Agent uses for data processing. For example, the Secure Agent uses the Data Integration Server to run data integration jobs and Process Server to run application integration and process orchestration jobs. Each Secure Agent service runs independently of the other services that run on the agent.

The independent services architecture provides the following benefits:

- The Secure Agent does not restart when you add a connector or package.
- Services are not impacted when another service restarts. For example, process orchestration jobs continue to run when the Data Integration Server restarts.
- Downtime during upgrades is minimized. The upgrade process installs a new version of the Secure Agent, updates connector packages, and applies configuration changes for the Data Integration Server. To minimize downtime, the old agent remains available and continues to run data integration jobs during the upgrade. The new version of the Secure Agent runs jobs that start after the upgrade process completes.

The services that run on a Secure Agent vary based on your licenses and the Informatica Intelligent Cloud Services that your organization uses.

The following table describes the services that can run on a Secure Agent and the Informatica Intelligent Cloud Services that use them:

Secure Agent service	Description	Used by...
API Microgateway Service	Manages Application Integration processes that run on the Secure Agent. **Preview Notice:** Effective in the Fall 2020 October release, API Microgateway Service is available for preview.	Application Integration, API Manager
B2B Processor	Runs B2B Gateway inbound and outbound process flows. **Note:** Do not configure this service unless you are instructed to do so by Informatica Global Customer Support.	B2B Gateway
CIH Processor	Runs Cloud Integration Hub publications and subscriptions for organizations that use a private publication repository.	Cloud Integration Hub
CMI Streaming Agent	Runs streaming ingestion jobs in the Mass Ingestion service.	Mass Ingestion service
Common Integration Components	Runs the shell scripts or batch commands in a Command Task step of a taskflow.	Data Integration
Database Ingestion	Runs database ingestion jobs in the Mass Ingestion service.	Mass Ingestion service

Secure Agent service	Description	Used by...
Data Integration Server	Runs data integration jobs such as mapping, task, and taskflow instances.	B2B Gateway, Cloud Integration Hub, Data Accelerator for Azure, Data Integration Data Profiling
EDC Search Agent	Discovers Enterprise Data Catalog data assets for Data Accelerator for Azure and for data catalog discovery in Data Integration. **Note:** Do not configure this service unless you are instructed to do so by Informatica Global Customer Support.	Data Accelerator for Azure, Data Integration
Elastic Server	Manages an elastic cluster and the elastic jobs that run on the cluster.	Data Integration
File Integration Service	Uses file transfer protocols such as HTTPS, AS2, and SFTP to receive files from a remote server or to send files to a remove server, or both.	B2B Gateway, Data Integration
Mass Ingestion	Runs file ingestion tasks and file listener jobs. **Note:** Do not configure this service unless you are instructed to do so by Informatica Global Customer Support.	Mass Ingestion service
OI Data Collector	Runs the data collectors and performs PowerCenter Integration Service grid auto-scaling for Operational Insights. **Note:** Do not configure this service unless you are instructed to do so by Informatica Global Customer Support.	Operational Insights
Process Server	Runs application integration processes, connectors, and connections.	Application Integration, Application Integration Console

Each service has a unique set of configuration properties, such as Tomcat and Tomcat JRE settings. You might need to configure a service or change the service properties to optimize performance or if you are instructed to do so by Informatica Global Customer Support. You configure a Secure Agent service independently from other services that run on the agent.

API Microgateway Service

The API Microgateway Service manages Application Integration processes that run on your organization's on-premises Secure Agent. Use the API Microgateway Service to expose managed APIs as API Microgateway proxies.

Preview Notice: Effective in the 2020 October release, API Microgateway Service is available for preview.

Preview functionality is supported for evaluation purposes but is unwarranted and is not production-ready. Informatica recommends that you use in non-production environments only. Informatica intends to include the preview functionality in an upcoming release for production use, but might choose not to in accordance

with changing market or technical circumstances. For more information, contact Informatica Global Customer Support. To use the functionality, your organization must have the appropriate licenses.

The API Microgateway Service supports the following API Manager access policies:

- IP filtering policy
- Rate limit policy
- Basic authentication

The API Microgateway Service provides REST APIs to create and deploy API Microgateway proxies. API consumers access the managed APIs deployed as API Microgateway proxies on your organization's on-premises environment. The Application Integration processes expose REST Service URL and SOAP Service URL endpoints.

Use the API Microgateway Service to build an API Microgateway proxy to an API endpoint to manage. The API Microgateway Service builds an API Microgateway as an immutable Docker image on your organization's Secure Agent machine. You then use the API Microgateway Service to deploy the Docker image in a container on the Secure Agent Docker runtime environment for API access. The API Microgateway applies the API access policies that you configure before forwarding the requests to Application Integration endpoints.

The following diagram shows the API Microgateway Service and API Microgateway components exposing a managed API in the on-premises environment:

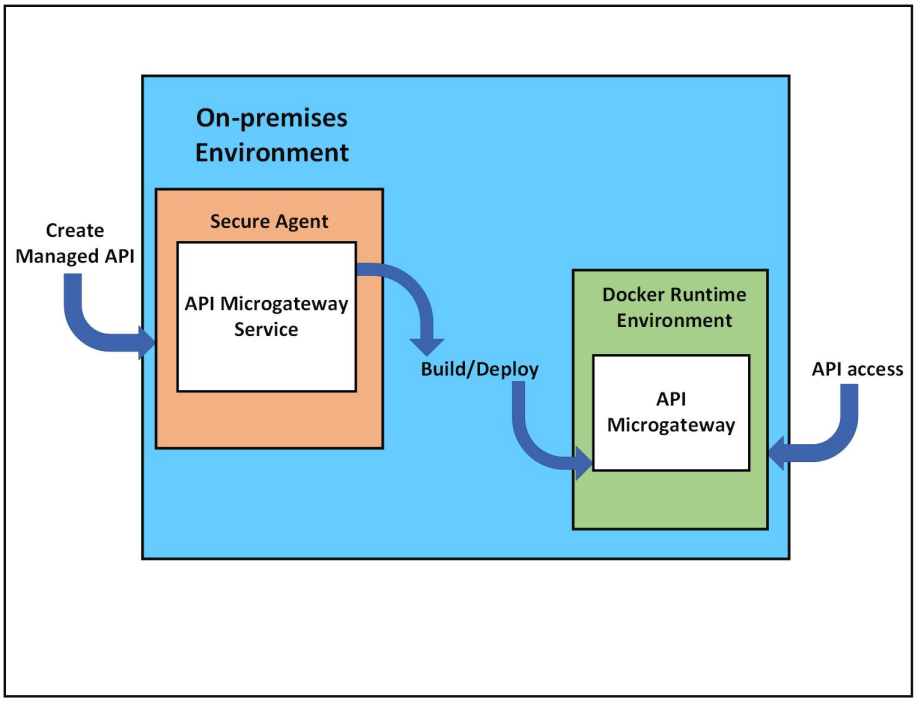

The Secure Agent Docker runtime environment hosts the Docker images using the blue-green deployment strategy to provide zero downtime during updates of the API Microgateway component.

Configuring API Microgateway Service

Configure the API Microgateway Service properties when you edit a Secure Agent in Administrator.

The following image shows the API Microgateway Service properties in the **System Configuration Details** area:

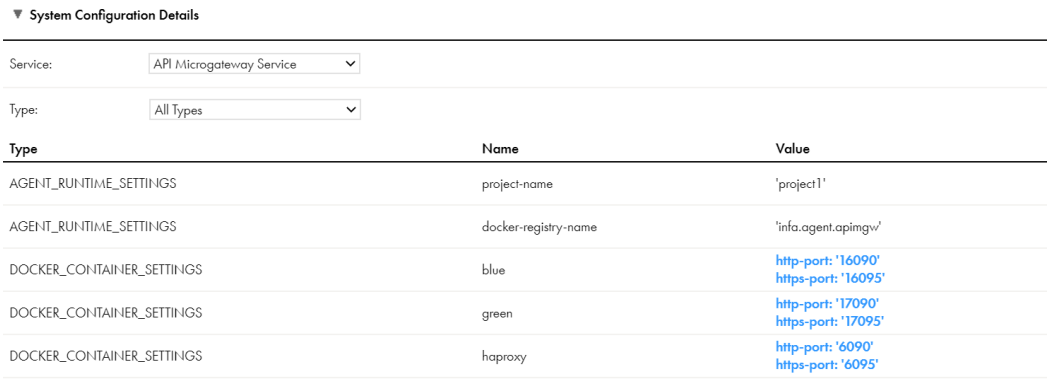

The following table describes the API Microgateway Service configurations:

Type	Name	Description
AGENT_RUNTIME_SETTINGS	project-name	Name of the project that stores the API configurations. You can change the name as per requirement, for example when you create a new project. **Note:** Project name must not contain the characters `/` or `/0`. If a project name includes restricted characters, the project creation fails.
AGENT_RUNTIME_SETTINGS	docker-registry-name	Name of the local Docker registry that contains all the named and tagged API Microgateway Docker images on the Secure Agent machine. **Note:** Docker image and tag names must not contain the following characters: `- _ , .` If a Docker image or tag name includes restricted characters, the image build fails.
DOCKER_CONTAINER_SETTINGS	blue	First Docker image container that deploys on the Secure Agent machine, alternates with green. You can change the following ports of the blue container: - http-port. The default value is: 16090 - https-port. The default value is: 16095
DOCKER_CONTAINER_SETTINGS	green	Second Docker image container that deploys on the Secure Agent machine, alternates with blue. You can change the following ports of the green container: - http-port. The default value is: 17090 - https-port. The default value is: 17095
DOCKER_CONTAINER_SETTINGS	haproxy	Router of the Docker image containers on the Secure Agent machine. Switches traffic between blue and green containers. You can change the following ports of the haproxy container: - http-port. The default value is: 6090 - https-port. The default value is: 6095

Note: To stop the API Microgateway, stop all three Docker image containers.

CMI Streaming Agent

Use the CMI Streaming Agent to define and deploy streaming ingestion tasks. You configure streaming ingestion tasks in the Mass Ingestion service.

A CMI Streaming Agent runs on an on-premise system and works in conjunction with the Mass Ingestion Streaming service. In an on-premise system, the CMI Streaming Agent runs the jobs deployed by Mass Ingestion Streaming. The agent provides status and statistics updates of each job.

Note: Prior to Spring 2020 April release of Informatica Intelligent Cloud Services Mass Ingestion service, CMI Streaming Agent was called Streaming Ingestion Agent.

CMI Streaming Agent properties

To change or optimize the behavior of the CMI Streaming Agent, configure agent properties for your run-time environment. Configure CMI Streaming Agent properties in the **System Configuration Details** area when you edit a Secure Agent.

You can configure Engine, Agent, and Script properties of a CMI Streaming Agent.
The following image shows some of the CMI Streaming Agent properties:

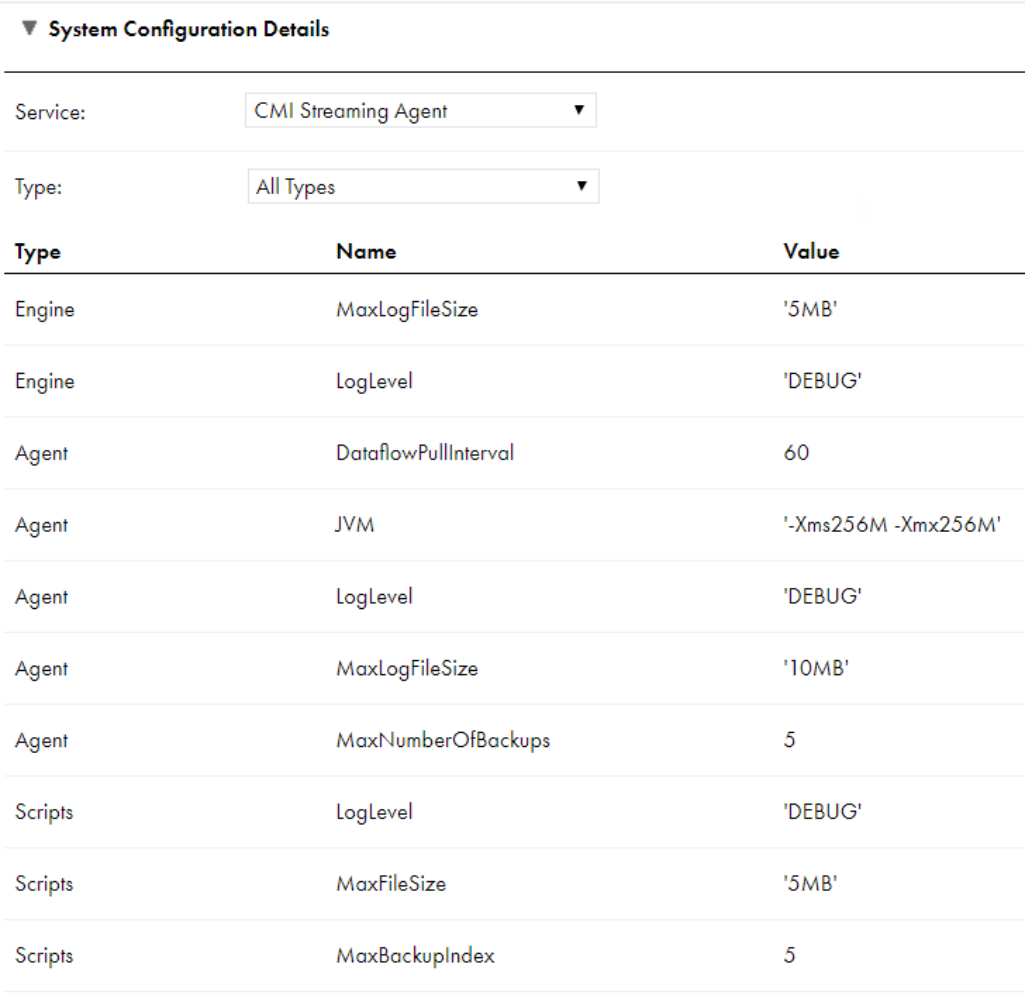

System Configuration Details

Service: CMI Streaming Agent

Type: All Types

Type	Name	Value
Engine	MaxLogFileSize	'5MB'
Engine	LogLevel	'DEBUG'
Agent	DataflowPullInterval	60
Agent	JVM	'-Xms256M -Xmx256M'
Agent	LogLevel	'DEBUG'
Agent	MaxLogFileSize	'10MB'
Agent	MaxNumberOfBackups	5
Scripts	LogLevel	'DEBUG'
Scripts	MaxFileSize	'5MB'
Scripts	MaxBackupIndex	5

You can configure the following CMI Streaming Agent properties:

Type	Property Name	Description
Engine	MaxLogFileSize	The maximum size of the log file that the engine can create. Default is 5 MB.
Engine	LogLevel	The log level for the engine.
Agent	DataflowPullInterval	The time interval after which the agent checks for updates in the task. Default is 60 seconds.
Agent	JVM	List of JVM properties for the agent. For example: `[-Xms256M -Xmx256M]`
Agent	LogLevel	The log level for the agent.
Agent	MaxLogFileSize	Maximum size of the log files that an agent can create. Default is 10 MB.
Agent	MaxNumberOfBackups	Maximum number of backup log files for the agent. Default is 5.
Scripts	LogLevel	The log level of the scripts.
Scripts	MaxFileSize	The maximum file size after which the log rolls over and creates a new file. Default is 10 MB.
Scripts	MaxBackupIndex	Maximum number of backup files maintained after rolling over. Default is 5.

Common Integration Components

The Common Integration Components service is the Secure Agent service that runs the commands specified in a Command Task step of a taskflow.

To view and configure the Common Integration Components service, you must have the license for the Common Integration Components service and command executor package.

You can optimize the performance of the Common Integration Components service by configuring some of its service properties. You can change service properties when you edit the Secure Agent.

All the requests that Common Integration Components service processes are logged in the following directory:

```
<Secure Agent installation directory>\apps\Common_Integration_Components\logs\<version>
```

You can view the log file for each command task in the following directory:

```
<Secure Agent installation directory>\apps\Common_Integration_Components\logs\command
\<Command_job ID>
```

Common Integration Components properties

To change or optimize the behavior of the Common Integration Components service, configure its properties in the **System Configuration Details** section when you edit a Secure Agent.

The following image shows some of the Common Integration Components service properties:

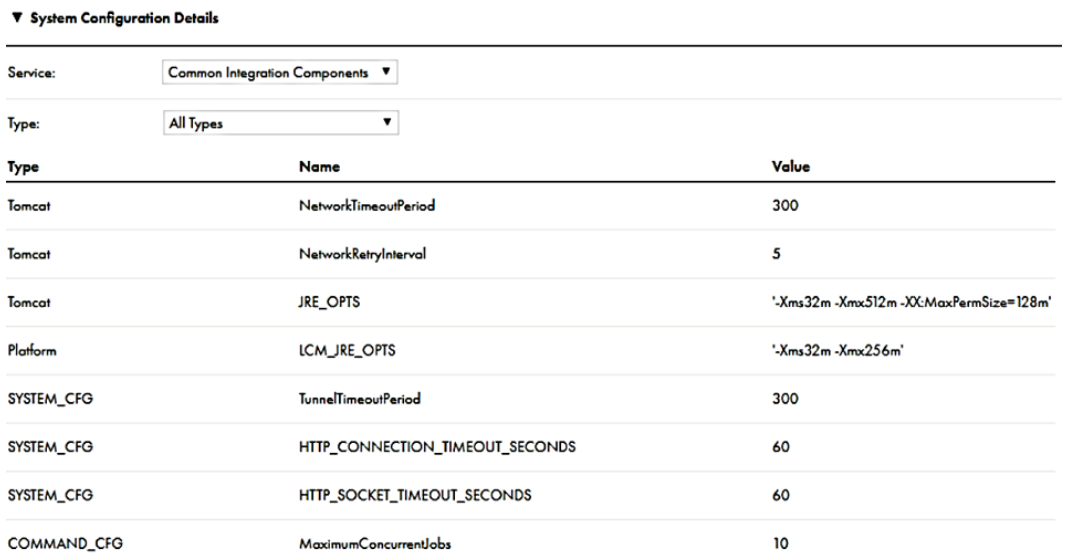

▼ System Configuration Details

Service: Common Integration Components

Type: All Types

Type	Name	Value
Tomcat	NetworkTimeoutPeriod	300
Tomcat	NetworkRetryInterval	5
Tomcat	JRE_OPTS	'-Xms32m -Xmx512m -XX:MaxPermSize=128m'
Platform	LCM_JRE_OPTS	'-Xms32m -Xmx256m'
SYSTEM_CFG	TunnelTimeoutPeriod	300
SYSTEM_CFG	HTTP_CONNECTION_TIMEOUT_SECONDS	60
SYSTEM_CFG	HTTP_SOCKET_TIMEOUT_SECONDS	60
COMMAND_CFG	MaximumConcurrentJobs	10

You can configure the following Common Integration Components service properties:

Type	Name	Description
Tomcat	JRE_OPTS	JRE VM options for the Apache Tomcat process.
Platform	LCM_JRE_OPTS	JRE options to start, stop, or get the status of the Apache Tomcat process. **Note:** Do not change the value of this property unless Informatica Global Customer Support instructs you to do so.
SYSTEM_CFG	HTTP_CONNECTION_TIMEOUT_SECONDS	The maximum amount of time, in seconds, that the Secure Agent waits to set up an HTTP connection to communicate with Informatica Intelligent Cloud Services. Default is 60. **Note:** Do not change the value of this property unless Informatica Global Customer Support instructs you to do so.
SYSTEM_CFG	HTTP_SOCKET_TIMEOUT_SECONDS	The maximum amount of idle time, in seconds, during the data packet transfer over an HTTP connection between the Secure Agent and Informatica Intelligent Cloud Services. Default is 60. **Note:** Do not change the value of this property unless Informatica Global Customer Support instructs you to do so.

Type	Name	Description
COMMAND_CFG	MaximumConcurrentJobs	The maximum number of concurrent command tasks that can be executed by a single Secure Agent. The default value is 10 for each Secure Agent in a Secure Agent group. For example, if there are 3 Secure Agents in a Secure Agent group, the maximum number of concurrent command tasks that the service can handle is 30. Any command execution requests beyond the maximum limit are queued and are executed when a Secure Agent is available.
Note: Do not change the values of other Common Integration Components service properties unless Informatica Global Customer Support instructs you to do so.		

Database Ingestion service

The Database Ingestion service (DBMI agent) enables you to define and run database ingestion tasks. You configure database ingestion tasks in the Mass Ingestion service.

After you download the Secure Agent to your runtime environment, the DBMI packages are pushed to the on-premises system where the Secure Agent runs, provided that you have custom licenses for both Mass Ingestion Databases and the DBMI packages. You can then optionally configure properties for the Database Ingestion service that runs on the Secure Agent.

Database Ingestion service properties

To change or optimize the behavior of the Database Ingestion service (DBMI agent) that the Secure Agent uses, configure Database Ingestion properties for your runtime environment.

To configure the properties, open your runtime environment. Under **System Configuration Details**, click **Edit**. Then select the **Database Ingestion** service and the **DBMI_AGENT_CONFIG** type.

The following table describes these agent properties:

Property	Description
maxTaskUnits	The maximum number of database ingestion tasks that can run concurrently on the on-prem machine where the Secure Agent runs. The default value is 10.
serviceLogRetentionPeriod	The number of days to retain each internal Database Ingestion service log file after the last update is written to the file. When this retention period elapses, the log file is deleted. The default value is 7 days. **Note:** Service logs are retained on the Secure Agent host where they are created: `<infaagent>/apps/Database_Ingestion/logs`.
taskLogRetentionPeriod	The number of days to retain each job log file after the last update is written to the file. When this retention period elapses, the log file is deleted. The default value is 7 days.

Property	Description
ociPath	For Oracle sources and targets, the path to the Oracle Call Interface (OCI) file oci.dll or libclntsh.so. For a DBMI agent that is running, this value is appended to the path that is specified in the PATH environment variable on Windows or in the LD_LIBRARY_PATH environment variable on Linux.
serviceUrl	The URL that the Database Ingestion service uses to connect to the Informatica Intelligent Cloud Services cloud.
logLevel	The level of detail to include in the logs that the Database Ingestion service produces. Options are: - TRACE - DEBUG - INFO - WARN - ERROR The default value is TRACE.
taskExecutionHeapSize	The maximum heap size, in gigabytes, for the Task Execution service. This value, in conjunction with maxTaskUnits property, affects the number of concurrent database ingestion tasks that can run on a Secure Agent. Try increasing the heap size to run more tasks concurrently. Enter this value followed by "g" for gigabytes, for example, '9g'. The default value is '8g'.

Database Ingestion Agent environment variables

To change or optimize the behavior of the Database Ingestion (DBMI) Agent, define the following environment variables:

Environment Variable	Description
DBMI_REPLACE_UNSUPPORTED_CHARS	For Microsoft Azure SQL Data Warehouse targets, controls whether a database ingestion job replaces characters in character data that the target cannot process correctly. To enable character replacement, set this environment variable to true. `DBMI_REPLACE_UNSUPPORTED_CHARS=true` Mass Ingestion Databases then uses the character that is specified in the DBMI_UNSUPPORTED_CHARS_REPLACEMENT environment variable to replace unsupported characters.
DBMI_UNSUPPORTED_CHARS_REPLACEMENT	If the DBMI_REPLACE_UNSUPPORTED_CHARS environment variable is set to true, specifies the character that replaces the characters in source data that a Microsoft Azure SQL Data Warehouse target cannot process correctly. Default value: ? (question mark)
DBMI_WRITER_CONN_POOL_SIZE	Indicates the number of connections that a database ingestion job uses to propagate the change data to the target. The default value is 8. Valid values are 4 through 8.

Environment Variable	Description
DBMI_WRITER_RETRIES_MAX_COUNT	If a network issue occurs while a database ingestion job is loading source data to an Amazon S3 or Microsoft Azure Data Lake Storage Gen2 target, indicates the maximum number of times that the database ingestion job retries a request to continue the initial load or incremental load. If all of the retries fail, the job fails. The default value is 5.
DBMI_WRITER_RETRIES_INTERVAL_IN_MILLIS	Specifies the time interval, in milliseconds, that a database ingestion job waits before retrying the request to continue the initial load or incremental load to an Amazon S3 or Microsoft Azure Data Lake Storage Gen2 target if a network issue occurs. The default value is 1000.

Note: After you define or change an environment variable, restart the Database Ingestion Agent for the changes to take effect.

Data Integration Server

The Data Integration Server is the Secure Agent service that runs data integration jobs such as mapping, task, and taskflow instances.

The Data Integration Server does not run elastic mappings and associated mapping tasks. If a taskflow contains a mapping task that is based on an elastic mapping, the Data Integration Server defers the mapping task to the Elastic Server and the Data Integration Server runs the other tasks in the taskflow.

You can optimize performance of the Data Integration Server by configuring some of its service properties. For example, you might want to change the network resiliency settings or the connection timeout period for the Secure Agent. You can change service properties when you edit the Secure Agent.

Data Integration Server resiliency

During temporary network issues, data integration tasks can continue to run while the Secure Agent tries to reestablish a connection. You can configure network resiliency properties for the Data Integration Server.

The following Data Integration Server properties determine how the Secure Agent tries to reestablish a connection:

NetworkTimeoutPeriod

Determines the length of time that the Secure Agent tries to reestablish communication with Informatica Intelligent Cloud Services. If communication is not established at the end of the time period, data integration tasks that were in progress stop running. The default value is 300 seconds.

NetworkRetryInterval

Determines the frequency with which the Secure Agent tries to contact Informatica Intelligent Cloud Services within the specified timeout period. The default value is five seconds.

For example, with the default settings, if the network is down, the Secure Agent tries to reestablish communication with Informatica Intelligent Cloud Services for 300 seconds. During the 300-second period, the Secure Agent tries to contact Informatica Intelligent Cloud Services every five seconds. If the Secure Agent reestablishes communication within the 300-second period, data integration tasks that are in progress

are not affected. If the Secure Agent is unable to reestablish communication within the 300-second period, the Secure Agent stops all data integration tasks that are in progress.

Data Integration Server properties

To change or optimize behavior of the Data Integration Server, configure the Data Integration Server properties. Configure Data Integration Server properties in the **System Configuration Details** area when you edit a Secure Agent.

The following image shows some of the Data Integration Server properties:

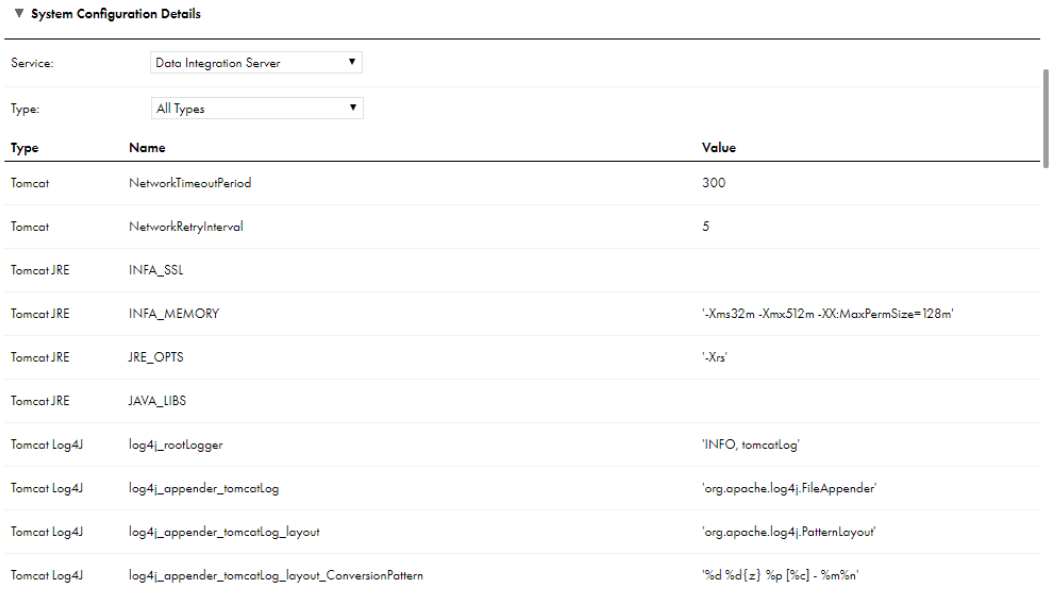

You can configure the following Data Integration Server properties:

Type	Name	Description
Tomcat	NetworkTimeoutPeriod	Amount of time, in seconds, that the Secure Agent tries to reestablish communication with Informatica Intelligent Cloud Services. Default is 300.
Tomcat	NetworkRetryInterval	Frequency, in seconds, in which the Secure Agent tries to contact Informatica Intelligent Cloud Services within the specified timeout period. Default is 5.
Tomcat JRE	JRE_OPTS	JRE VM options for the Apache Tomcat process.
Tomcat JRE	INFA_MEMORY	JRE VM options that are set for virtual machine memory for the Apache Tomcat process.
DTM	AgentConnectionTimeout	Number of seconds that the Secure Agent communication requests to wait before it times out. Default is 5.

Type	Name	Description
DTM	JVMOpton1 - JVMOption5	JVM options that configure advanced properties for the Data Integration Server such as the maximum and minimum JVM heap size, the maximum record size for Intelligent Structure Discovery, or proxy settings for certain connectors. For example, to change the maximum JVM heap size from the default value of 512 MB to 2048 MB, you might set JVMOption1 to `'-Xmx2048m'`. By default, you can configure up to five advanced properties using JVMOption1 through JVMOption5. To configure additional properties, you can add custom DTM properties for the Data Integration Server named JVMOption6, JVMOption7, etc. Ensure that the option numbers are sequential and that you do not skip numbers. For information about the JVM options that you can set, see the Data Integration help, the help for the appropriate connector, or the Knowledge Base on Informatica Network.
Note: Do not change the values of other Data Integration Server properties unless Informatica Global Customer Supports instructs you to do so.		

Elastic Server

The Elastic Server is the Secure Agent service that manages an elastic cluster and elastic jobs that run on the cluster.

You can configure the service properties to specify the level of detail that the Elastic Server writes to log files. There are also service properties that you must configure to set up the kops and Secure Agent roles. For more information, see *Data Integration Elastic Administration*.

Elastic Server properties

To change the behavior of the Elastic Server, configure the Elastic Server properties in the **System Configuration Details** area when you edit a Secure Agent.

The following image shows the Elastic Server properties:

You can configure the following Elastic Server properties:

Type	Name	Description
LOG4J_CFG	log4j_app_log_level	Level of detail that the Elastic Server writes to log files. Enter the logging level as a string, such as 'INFO.' As the logging level increases, the messages that the Elastic Server writes to log files include the messages in the prior logging levels. For example, if the logging level is INFO, the log contains FATAL, ERROR, WARNING, and INFO code messages. The following values are valid: 1. FATAL. Includes nonrecoverable system failures that cause the service to shut down or become unavailable. 2. ERROR. Includes connection failures, failures to save or retrieve metadata, and service errors. 3. WARNING. Includes recoverable system failures or warnings. 4. INFO. Includes system and service change messages. 5. TRACE. Logs user request failures. 6. DEBUG. Logs user request logs.
AWS_CFG	agent_role_external_id_key	External ID that the Secure Agent specifies when the agent attempts to assume the kops role. Required if you configure an external ID in the trust relationship of the kops role. Ignore this property in an Azure environment.
AWS_CFG	privileged_role_arn_key	ARN of the kops role. Required when you set up separate kops and Secure Agent roles in an AWS environment. Ignore this property in an Azure environment.
AWS_CFG	role_session_duration_secs_key	Session duration of the AWS AssumeRole API in seconds. By default, the session duration is 1800 seconds (30 minutes). Overrides the maximum CLI/API session duration that is configured for the kops role. If the session duration configured for the Elastic Server is longer than session duration for the kops role, the Secure Agent might fail to assume the kops role. Ignore this property in an Azure environment.

File Integration Service

Use the File Integration Service to transfer files between your organization and remote file servers.

The File Integration Service is a Secure Agent service that the agent uses to run advanced file transfer protocols such as AS2.

You must configure file servers before your organization can receive files from remote partners. Configure your organization's file server associated with the File Integration Service on the File Servers page in Administrator. Configuration includes properties such as file server details, encryption methods, and allowed file types.

To stop or start the File Integration Service, you stop or start the file server that uses the service.

For information about configuring file servers, see "File server configuration process" on page 164.

To use the File Integration Service, your organization must have the appropriate license. To configure the File Integration Service, you must be assigned the Admin role.

Mass Ingestion (Files)

To change or optimize the behavior of Mass Ingestion Files, configure the Mass Ingestion properties. Configure the properties in the **System Configuration Details** area when you edit a Secure Agent.

You can configure the following properties:

Type	Name	Description
AGENT_RUNTIME_SETTINGS	file-listener-snapshot-dir	The directory where the snapshots of a new file listener components are added. You can add the following directory paths: - A path relative to the `MassIngestionRuntime` directory. For example, `../data/monitor`. - The absolute path. For example, `<Secure agent installation directory>/apps/MassIngestionRuntime/data/monitor` where *Secure agent installation directory* is the name of the directory where the secure agent is installed. **Note:** Use the snapshot directory shared with all agents when multiple Secure Agents are present in a group.
AGENT_RUNTIME_SETTINGS	mi-task-workspace-dir	The directory to a custom location in the agent. A custom directory in the agent that file ingestion tasks use as an intermediate staging area when transferring files to a target. The path can be a shared location, mounted location, or a location apart from the default location in the agent.
AGENT_RUNTIME_SETTINGS	file-listener-max-pool-size	The maximum number of threads to execute the file listener. Default is 20.
AGENT_RUNTIME_SETTINGS	file-listener-core-pool-size	The total number of threads. Default is 20.
AGENT_RUNTIME_SETTINGS	ftp-receive-socket-buffer-size	The buffer size for FTP inbound packets. Default is 16 bytes.
AGENT_RUNTIME_SETTINGS	ftp-send-socket-buffer-size	The buffer size for FTP outbound packets. Default is 16 bytes.
AGENT_RUNTIME_SETTINGS	http-client-timeout	The timeout duration in seconds for Agent requests to Informatica Intelligent Cloud Services. Default is 30 seconds.

Type	Name	Description
PGP_SETTINGS	public-keyring-path	The directory to store the public key ring. You can add the following directory paths: - A path relative to the directory where mass ingestion is installed. For example, `../data/pubring.pkr` where *pubring.pkr* is the name of the file where you store the public key ring. - The absolute path. For example, `<Secure agent installation directory>/apps/MassIngestionRuntime/data/pubring.pkr` where *pubring.pkr* is the name of the file where you store the public key ring and *Secure agent installation directory* is the name of the directory where the agent is installed.
PGP_SETTINGS	secret-keyring-path	The directory to store the secret key ring. You can add the following directory paths: - A path relative to the directory where mass ingestion is installed. For example, `../data/secring.pkr` where *secring.pkr* is the name of the file where you store the secret key ring. - The absolute path. For example, `<Secure agent installation directory>/apps/MassIngestionRuntime/data/secring.pkr` where *secring.pkr* is the name of the file where you store the secret key ring and *Secure Agent installation directory* is the name of the directory where the agent is installed.
JVM_SETTINGS	app-heap-size	The minimum and maximum heap sizes of the Mass Ingestion Files application. Default is -Xms256m -Xmx2048m.
JVM_SETTINGS	lcm-heap-size	The minimum and maximum heap sizes of life-cycle management scripts. Default is -Xms32m -Xmx128m.

Process Server

Process Server is the Secure Agent service that executes Application Integration processes, connectors, and connections.

When you deploy Application Integration assets to the Secure Agent, you deploy them to Process Server. When you run an asset, Process Server executes it.

The PostgreSQL database comes with the Process Server service of the Secure Agent and stores the metadata that Process Server collects and generates.

Find the PostgreSQL directory at the following location on your system:

```
<Secure Agent installation directory>\apps\process-engine\data\PostGreSql
```

Process Server properties

To change or optimize the behavior of Process Server, configure Process Server properties. You can configure the server, Secure Agent group, Java Virtual Machine, connector, database, and logging properties.

The following image shows some Process Server properties:

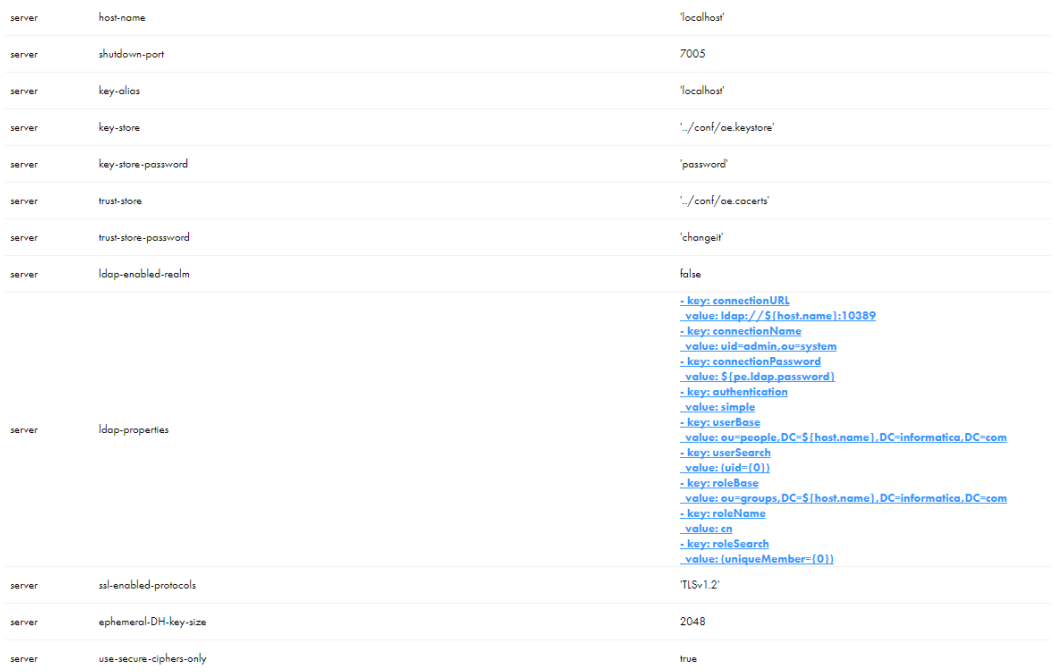

server	host-name	'localhost'
server	shutdown-port	7005
server	key-alias	'localhost'
server	key-store	'../conf/ae.keystore'
server	key-store-password	'password'
server	trust-store	'../conf/ae.cacerts'
server	trust-store-password	'changeit'
server	ldap-enabled-realm	false
server	ldap-properties	- key: connectionURL value: ldap://${host.name}:10389 - key: connectionName value: uid=admin,ou=system - key: connectionPassword value: ${pe.ldap.password} - key: authentication value: simple - key: userBase value: ou=people,DC=${host.name},DC=informatica,DC=com - key: userSearch value: (uid={0}) - key: roleBase value: ou=groups,DC=${host.name},DC=informatica,DC=com - key: roleName value: cn - key: roleSearch value: (uniqueMember={0})
server	ssl-enabled-protocols	'TLSv1.2'
server	ephemeral-DH-key-size	2048
server	use-secure-ciphers-only	true

You can configure the following server properties:

Name	Communication Method	Description
`host-name`	Secure Agent Channel	The host name of the Process Engine server.
`shutdown-port`	Secure Agent Channel	Process Server Tomcat shutdown port.
`key-alias`	HTTPS	The identifier of the keystore record that contains security keys for HTTPS communication.

`key-store`	HTTPS	The path and file name of the key store file that Application Integration uses for HTTPS communication.
		When you install the Secure Agent, you can find the key store in the following default location:
		`<Secure Agent installation directory>/apps/process-engine/conf/ae.keystore`
		You can also enter a relative path. For example, if the current working directory is the Secure Agent installation directory, enter the following value to point to the `ae.keystore` file:
		`../conf/ae.keystore`
		Note: The file path can contain only forward slashes (`/`).

Name	Communication Method	Description
key-store-password	HTTPS	The key store password. The default password is password.
trust-store	HTTPS	The path and file name of the trust store file that Application Integration uses for HTTPS communication. When you install the Secure Agent, you can find the trust store in the default location: <Secure Agent installation directory>/apps/process-engine/conf/ae.cacerts You can also enter a relative path. For example, if the current working directory is the Secure Agent installation directory, enter the following value to point to the ae.cacerts file: ../conf/ae.cacerts **Note:** The file path can contain only forward slashes (/). If you want to import public certificates for service endpoint authentication, place them in the following location: <Secure Agent installation directory>/apps/process-engine/conf/certs
trust-store-password	HTTPS	The trust store password. The default password is changeit. You can change the password.
ldap-enabled-realm	HTTP/HTTPS	Set this property to True if you want to use an LDAP provider for authentication. Use the LDAP provider as a centralized form of authentication when you have clustered Secure Agents.
ldap-properties	HTTP/HTTPS	The LDAP properties that you need to configure. Edit the existing properties to suit your LDAP provider. **Note:** Your LDAP password does not appear on screen. The value of $(pe.ldap.password) is taken from the PE_LDAP_PASSWORD environment variable
ssl-enabled-protocols	HTTPS	The TLS protocol to use. The default protocol, TLSv1.2, is the most secure protocol. Change this value to an older version like TLSv1.0 or TLSv1.1 only if you face compatibility issues
ephemeral-DH-key-size	HTTPS	The key length of the secure algorithm. The default value is 2048. Change this value only if you face compatibility issues.
use-secure-ciphers-only	HTTPS	Limits the set of ciphers used during a call to the endpoint to secure ciphers only. The default value is True. Change this value to false only if you face compatibility issues.

You can configure the following Secure Agent group ('cluster' on the UI) properties:

Name	Communication Method	Description
`name`	HTTP/HTTPS	The name of the Secure Agent group.
`primary-node`	HTTP/HTTPS	Set this property to `true` if you want the Secure Agent to be the master agent. When you select a master agent, you create a Secure Agent cluster. In a cluster, all Secure Agents share the postgreSQL database of the master Secure Agent.
`load-balance-url`	HTTP/HTTPS	The load balancer URL that you can use to invoke the process deployed to the Secure Agent. Applicable if you have a load balancer.

You can configure the following Java Virtual Machine properties:

Name	Communication Method	Description
`min-heap`	Secure Agent Channel	The minimum heap memory that Process Server allocates to the Tomcat JVM.
`max-heap`	Secure Agent Channel	The maximum heap memory that Process Server allocates to the Tomcat JVM.
`additional-properties`	Secure Agent Channel	A custom system property that you can add to the Tomcat JVM set. For example, you can set the custom property –`Dsun.net.inetaddr.ttl=60`

You can configure the following connector properties:

Name	Communication Method	Description
`http-port`	HTTP	The HTTP port to which the Secure Agent sends data. The default port is 7080. For more information about the construction of REST and SOAP endpoint URLs, see the Application Integration help.
`http-maxThreads`	HTTP	The maximum number of connections that Process Server creates with Application Integration over HTTP.
`http-connectionTimeout`	HTTP	The maximum time, in milliseconds, that Process Server waits for an HTTP connection to reply.
`https-port`	HTTPS	The HTTPS port to which the Secure Agent sends data. The default port is 7443. For more information about the construction of REST and SOAP endpoint URLs, see the Application Integration help.
`https-maxThreads`	HTTPS	The maximum number of connections that Process Server creates with Application Integration over HTTPS.

Name	Communication Method	Description
`https-connectionTimeout`	HTTPS	The maximum time, in milliseconds, that Process Server waits for an HTTPS connection to reply.
`secure-channel maxThreads`	Secure Agent Channel	The maximum number of connections that Process Server creates with Application Integration.
`secure-channel-connectionTimeout`	Secure Agent Channel	The maximum time, in milliseconds, that Process Server waits for a connection to reply.

You can configure the following database properties:

Name	Communication Method	Description
`type`	Secure Agent Channel	The database type that Process Server runs on. **Important:** Do not change this setting. The Application Integration Secure Agent does not support other databases.
`driver`	Secure Agent Channel	The database driver that Process Server runs on. **Important:** Do not change this setting. The Informatica Cloud Secure Agent does not support other databases.
`URL`	Secure Agent Channel	URL at which Process Server accesses the database. **Important:** Do not change this setting. The Informatica Cloud Secure Agent does not support other databases.
`maxActive`	Secure Agent Channel	The maximum number of active connections allocated to Process Server database at the same time.
`maxIdle`	Secure Agent Channel	The maximum number of connections that can remain idle at a time in the database. Process Server releases connections if the number of idle connections crosses this number.
`maxWait`	Secure Agent Channel	The maximum time that the database waits for a connection if none are available.
`connection-properties`	Secure Agent Channel	Key-value pairs of database connection properties. Some keys are available by default. Do not delete the default keys. You can, however, change the values of these keys. You can add other key-value pairs. For example, you can add the following key-value pair: key: `autoReconnect` value: `true`

You can configure the following logging properties:

Name	Communication Method	Description
1catalina_org_apache_juli_FileHandler_level	Secure Agent Channel	The level of logging in the file: < Secure Agent installation directory> \apps\process-engine\logs \catalina.log. Default: FINE
2localhost_org_apache_juli_FileHandler_level	Secure Agent Channel	The level of logging in the file: < Secure Agent installation directory> \apps\process-engine\logs \localhost.log. Default: FINE
3manager_org_apache_juli_FileHandler_level	Secure Agent Channel	The level of logging in the file: < Secure Agent installation directory> \apps\process-engine\logs \manager.log. Default: FINE
4host_manager_org_apache_juli_FileHandler_level	Secure Agent Channel	The level of logging in the file: < Secure Agent installation directory> \apps\process-engine\logs \host-manager.log. Default: FINE
java_util_logging_ConsoleHandler_level	Secure Agent Channel	The level of logging in the **CMD** window that appears when you start Tomcat. Default: FINE

Name	Communication Method	Description
org_apache_catalina_core_ContainerBase_Catalina_localhost_level	Secure Agent Channel	The level of logging in the localhost.log file when you host Tomcat on a virtual machine. Default: INFO
org_apache_catalina_core_ContainerBase_Catalina_localhost_manager_level	Secure Agent Channel	The level of logging in the manager.log file when you host Tomcat on a virtual machine. Default: INFO
org_apache_catalina_core_ContainerBase_Catalina_localhost_host-manager_level	Secure Agent Channel	The level of logging in the host-manager.log file when you host Tomcat on a virtual machine. Default: INFO

Default connection database properties

The following table describes the default keys that are available for the connection-properties database property:

Key	Description
timeBetweenEvictionRuns	The number of milliseconds that Process Server waits in-between runs of the idle object evictor thread.
testOnBorrow value	Process Server validates objects before borrowing objects from the pool. If Process Server cannot validate the object, it drops the object from the pool. Then, Process Server tries to borrow another object.
testWhileIdle	Process Server validates objects by the idle object evictor (if one exists). If Process Server cannot validate the object, it drops the object from the pool.
validationQuery value	The SQL query that validates connections from this pool before returning them to the caller. If you specify this property, the query must be an SQL SELECT statement that returns at least one row.

Logging levels

The following table describes the levels that you can configure for Process Server **logging** properties:

Level	Description
SEVERE	Logs errors.
WARNING	Logs potentially harmful situations.
INFO	Logs informational events that show the high-level progress of the application.
CONFIG	Logs informational events in more detail than at the INFO level.
FINE	Logs fine-grained informational events that you can use to debug an application.
FINER	Logs fine-grained informational events in more detail than at the FINE level.
FINEST	Logs all events.

Process Server sizing recommendations

Configure the Process Server service of the Secure Agent according to your workload.

Use the following sizing recommendations to optimize resources:

Recommendation	Small	Medium	Large
Process Count	75	175	350
Resource Cache (MB)	75	175	350
Work Manager Thread Pool Min	50	100	150
Work Manager Thread Pool Max	250	500	750
JVM Min Heap (MB)	Default	768	1024
JVM Max Heap (MB)	Default	Default	4096

The default JVM Min Heap is 512 MB, and the default JVM Max Heap is 1536 MB.

To configure the Process Count, Resources Cache, Work Manager Thread Pool Min, and Work Manager Thread Pool Max, go to the **Server Configuration** section of the Application Integration Console service.

When you start the Process Server on UNIX operating systems, you might receive the following error:

```
Cannot write to temp location [/tmp]
```

This error occurs because UNIX limits the number of files that can be created by a single process. The maximum number of files that can be created by a single process is 1024.

To avoid this error, Informatica recommends that you increase the open files limit to atleast 10 times the default value of 1024. Contact your system administrator to increase the value of any other relevant parameters such as max user processes.

For more information about Process Server sizing for a Secure Agent, see the following document:

https://network.informatica.com/docs/DOC-17439

Communication with the Secure Agent

Informatica Intelligent Cloud Services sends data from the Secure Agent to the Process Server through the Secure Agent Channel or a through a direct HTTP or HTTPS link.

The Secure Agent communicates with Process Server in two ways:

The Secure Agent Channel

A secure channel that creates tunnels between each connected Secure Agent and Process Server.

HTTP or HTTPS

A protocol over which the Secure Agent directly sends data to Process Server. If you use this communication method, Informatica Intelligent Cloud Services validates your credentials against the authentication provider.

For more information on the communication method that each Process Server property uses, see the Process Server Propertieson page 129.

The following image shows the two methods of communication between the Secure Agent and Process Server:

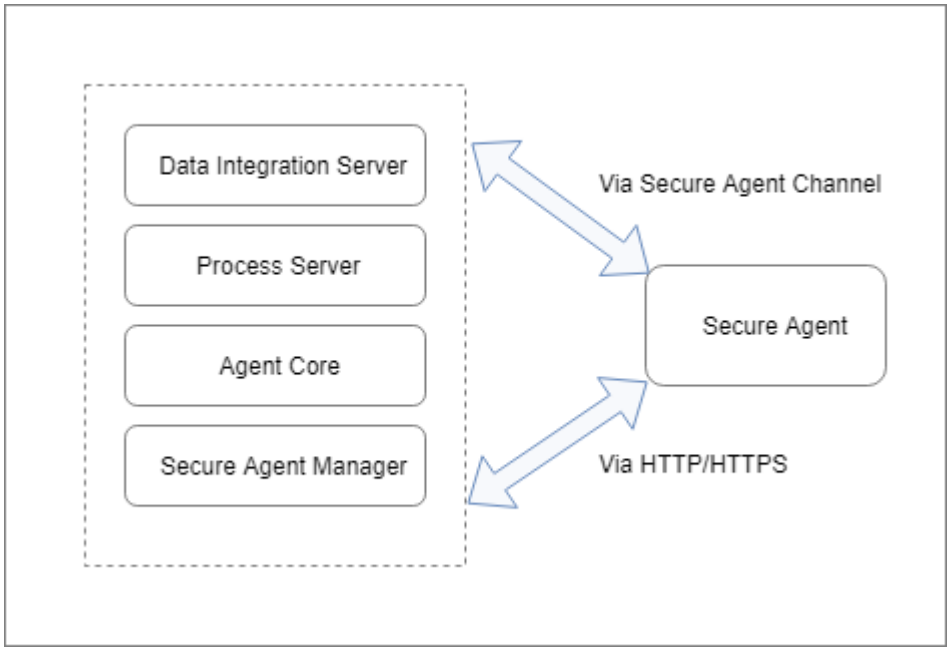

Secure Agent configurations for Process Server

Deploy an asset to a single Secure Agent, a Secure Agent group, or a Secure Agent cluster based on your business needs.

When you deploy Application Integration processes, connections, or service connectors to a Secure Agent, you deploy the assets to the Process Server service of the Secure Agent. All Secure Agents with the Process Server service use a PostgreSQL database.

You can deploy an asset to the following Secure Agent configurations:

Single Secure Agent

A single Secure Agent might be the only agent in a group, or one of many agents on a group.

For more information, see Deploy to a Single Secure Agenton page 137.

Secure Agent group

A Secure Agent group contain multiple agents. When you deploy an asset to a Secure Agent group, Informatica Intelligent Cloud Services performs load balancing. Use the Secure Agent group configuration to distribute requests if you process stateless requests or use the Secure Agent only to serve OData requests.

For more information, see Deploy to a Secure Agent Groupon page 138.

Secure Agent Cluster

A Secure Agent Cluster is an agent group that has a master securer agent. Use the Secure Agent Cluster configuration when you want all Process Servers to receive information about process execution activity.

For more information, see Deploy to a Secure Agent Clusteron page 138.

The following table summarizes the process execution of the Secure Agent in different scenarios:

	Single Secure Agent	Secure Agent Group	Secure Agent Cluster
Agent Available	Process executes.	Process executes.	Process executes.
Agent Unavailable	Process does not execute.	Process executes on any available Secure Agent.	Process executes on any available Secure Agent.
Agent Stops Mid-Execution	Process does not execute.	Process stops when the Secure Agent stops.	Process continues to execute on another Secure Agent.

Deploy to a single Secure Agent

You can deploy an asset directly to one agent in a group.

When you deploy an asset to a single Secure Agent, no other Process Server in the Secure Agent group receives the asset definitions.

The following image shows a sample configuration where process X is deployed directly to Secure Agent 1:

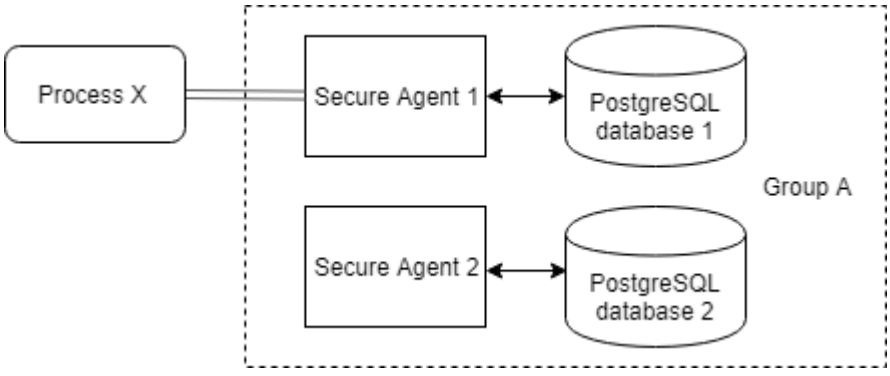

Only Secure Agent 1 can execute Process X. If Secure Agent 1 is unavailable, the process does not execute.

Deploy to a Secure Agent group

A Secure Agent group contains multiple agents. You can deploy an asset to a Secure Agent group.

Use the Secure Agent group configuration to distribute requests if you process stateless requests or use the Secure Agent only to serve OData requests.

In a Secure Agent group, Informatica Intelligent Cloud Services dispatches incoming requests to available Secure Agents in a round-robin manner.

When you deploy an asset to a Secure Agent group, you use a load balanced configuration. Informatica Intelligent Cloud Services performs the load balancing. You can also use a custom load balancer by setting the `load-balance-url` Process Server property. For more information, see Process Server Properties on page 129.

For more information about Secure Agent groups, see Secure Agent Groups with Multiple Agents on page 86.

All Secure Agents in a group use individual PostgreSQL databases. When you deploy an asset to a Secure Agent group, all Process Servers within the group receive details about new or updated asset definitions. However, the other Process Servers in the group do not receive details about the execution activity of an asset. For example, if a Secure Agent within the group fails during process execution, the process does not continue to execute on another Secure Agent within the group.

The following image shows a sample configuration where process X is deployed to Secure Agent group A:

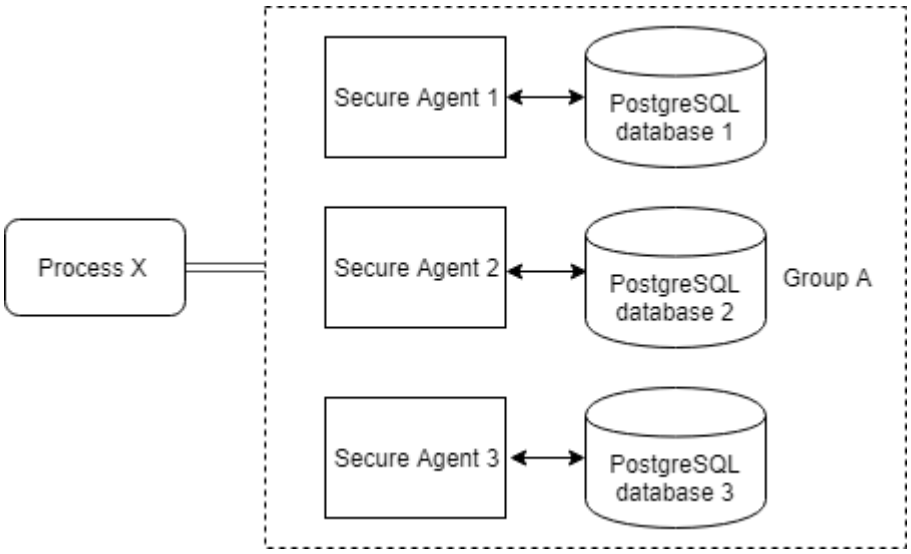

If you modify and re-publish process X, all three Secure Agents receive the updated definition. Any Secure Agent can execute the process.

For example, if the process is invoked and Secure Agent 1 and Secure Agent 2 are unavailable, the load balanced configuration ensures that Secure Agent 3 executes process X. However, Secure Agent 1 and Secure Agent 2 do not receive information about whether the process has faulted or completed successfully. If Secure Agent 3 stops while executing the process X, the process does not execute further.

Deploy to a Secure Agent Cluster

A Secure Agent Cluster is an agent group with a master Secure Agent. You can deploy an asset to a Secure Agent Cluster.

When you deploy an asset to a Secure Agent Cluster, all Process Servers receive information about process execution activity. The master Secure Agent receives information and informs the other Secure Agents. If a

Secure Agent fails during process execution, the process continues to executes on another Secure Agent within the cluster.

All Process Servers in a cluster share the PostgreSQL database of the master agent.

To define the master Secure Agent, use the `primary-node` Process Server property. For more information, see Process Server Properties on page 129.

The following image shows a sample configuration where process X is deployed to Secure Agent Cluster A:

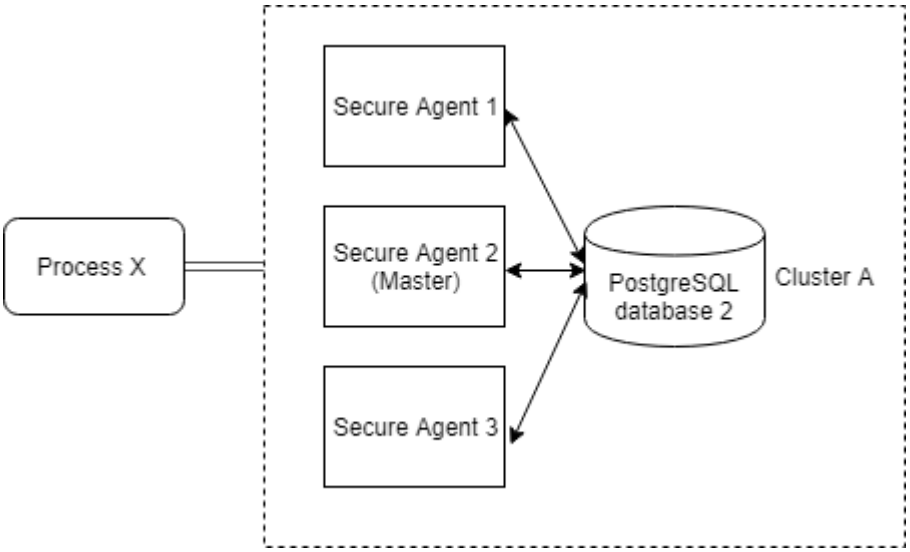

If Secure Agent 3 starts executing process X but stops midway, either Secure Agent 1 or Secure Agent 2 continues to execute the process.

Managing the PostgreSQL database on Windows

Use binaries and utility scripts to manage the PostgreSQL database.

Important: To manage the PostgreSQL database, you must log in as a user who does not have system administrator rights. A system administrator will be unable to run PostrgreSQL binaries and utility scripts.

Informatica has created some utility scripts based on PostgreSQL binaries. These utility scripts make it easier for you to manage the PostgreSQL database.

The following directories contain files for the PostgreSQL database:

- PostgreSQL binaries: `<Secure Agent installation directory>\apps\process-engine\data\db\postgresql-windows-x64-binaries\pgsql\bin`
- PostgreSQL utility scripts: `<Secure Agent installation directory>\apps\process-engine\data\db\util`
- PostgreSQL logs: `<Secure Agent installation directory>\apps\process engine\logs\PostGreSql\postgresql.log`
- PostgreSQL data: `<Secure Agent installation directory>\apps\process engine\data\PostGreSql\Data`

For more information about PostgreSQL scripts, see the PostgreSQL help at https://www.postgresql.org/docs/current/static/index.html.

Some of the following sections contain sample commands that use these default values:

- Default database name: activevos
- Default database user name: bpeluser
- Default database password: bpel

Backing up the PostgreSQL database on Windows

Use the script `db_backup.bat` to back up the PostgreSQL database.

To back up the PostgreSQL database, perform the following steps:

1. Go to `<Secure Agent installation directory>\apps\process-engine\data\db\util`.
2. Run the following command:
   `db_backup.bat <dbusername> <dbpassword> <path to backup file along with name of backup file with a ".dump" extension>`.

For example, the Secure Agent creates the backup file '`BackupFile1.dump`' in the location `C:\postgre\backup` if you run the following command:

```
db_backup.bat bpeluser bpel "C:\postgre\backup\BackupFile1.dump",
.
```

Restoring the PostgreSQL database on Windows

Use the command `db_restore.bat` to restore the PostgreSQL database from a backup file.

To restore the PostgreSQL database file from a backup file, perform the following steps:

1. Go to `<Secure Agent installation directory>\apps\process-engine\data\db\util`.
2. Run the following command:
   `db_restore.bat <dbusername> <dbpassword> <path to dump file>`.

For example, you use the file `BackupFile1.dump` to restore the PostgreSQL database if you run the following command:

```
db_restore.bat bpeluser bpel "C:\postgre\backup\BackupFile1.dump",
```

Resetting the PostgreSQL database on Windows

Shut down the PostgreSQL database and then use the command `db_reset.bat` to reset it.

To reset the PostgreSQL database to its original state, perform the following steps:

1. Go to `<Secure Agent installation directory>\apps\process-engine\data\db\util`
2. To shut down the server, run the following command:
   `server_stop.bat`
3. To reset the PostgreSQL database, run the following command:
   `db_reset.bat`

Starting the PostgreSQL server on Windows

To start the Process server on Windows, perform the following steps:

1. Go to `<Secure Agent installation directory>\apps\process-engine\data\db\util`.

2. Run the following script:
   `server_start.bat`.

Stopping the PostgreSQL server on Windows

To stop the PostgreSQL server on Windows, perform the following steps:

1. Go to `<Secure Agent installation directory>\apps\process-engine\data\db\util`.
2. Run the following script:
   `server_stop.bat`.

Getting the PostgreSQL server status on Windows

To get the status of the PostgreSQL server on Windows, perform the following steps:

1. Go to `<Secure Agent installation directory>\apps\process-engine\data\db\util`.
2. Run the following script:
   `server_status.bat`.

Vacuuming the PostgreSQL database on Windows

Vacuum the PostgreSQL database to delete obsolete tuples and gain space. You use the script `db_maintenance.bat` to vacuum the PostgreSQL database.

By default, the PostgreSQL database autovacuums. If you want to manually vacuum the database, perform the following steps:

1. Go to `<Secure Agent installation directory>\apps\process-engine\data\db\util`.
2. To vacuum the entire database, run the following command:
   `db_maintenance.bat <dbusername> <dbpassword> vacuum`
3. To vacuum a single table, run the following command:
   `db_maintenance.bat <dbusername> <dbpassword> vacuum <tablename>`

For example, you vacuum the 'aeprocesslogdata' table if you run the following command:

```
db_maintenance.bat bpeluser bpel vacuum aeprocesslogdata
```

Reindexing the PostgreSQL database on Windows

Use the reindexing option to clean the index and free up space after you vacuum data on PostgreSQL. You use the script `db_maintenance.bat` to reindex the PostgreSQL database.

To reindex the PostgreSQL database, perform the following steps:

1. Go to `< Secure Agent installation directory>\apps\process-engine\data\db\util`.
2. To reindex the entire database, run the following command:
   `db_maintenance.bat <dbusername> <dbpassword> reindex`
3. To reindex a single table, run the following command:
   `db_maintenance.bat <dbusername> <dbpassword> reindex <tablename>`

For example, you reindex the 'aeprocesslogdata' table if you run the following command:

```
db_maintenance.bat bpeluser bpel reindex aeprocesslogdata
```

Resetting transaction logs on Windows

If the PostgreSQL server does not start because of corruption of the control information, use the command `pg_resetxlog.exe` to reset the control information.

To reset the PostgreSQL database control information, perform the following steps:

1. Go to `<Secure Agent installation directory>\apps\process-engine\data\db\postgresql-windows-x64-binaries\pgsql\bin`.
2. Run the following command:
   `pg_resetxlog.exe -D <path to postgreSQL data directory>`

For example, you reset transactions logs in the 'Data' directory if you run the following command:

```
pg_resetxlog.exe -D "C:\postgre\apps\process-engine\data\PostGreSql\Data"
```

Managing the PostgreSQL database on Linux

Use binaries utility scripts to manage the PostgreSQL database.

Important: To manage the PostgreSQL database, you must log in as a user who does not have root access. A root user will be unable to run PostrgreSQL binaries and utility scripts.

Informatica has created some utility script based on PostgreSQL binaries. These utility scripts make it easier for you to manage the PostgreSQL database.

The following directories contain files for the PostgreSQL database:

- PostgreSQL binaries: `<Secure Agent installation directory>/apps/process-engine/data/db/postgresql-linux-x64-binaries/pgsql/bin`
- PostgreSQL utility scripts: `<Secure Agent installation directory>/apps/process-engine/data/db/util`
- PostgreSQL logs: `<Secure Agent installation directory>/apps/process-engine/logs/PostGreSql/postgresql.log`
- PostgreSQL data: `<Secure Agent installation directory>/apps/process-engine/data/PostGreSql/Data`

For more information about PostgreSQL scripts, see the PostgreSQL help at https://www.postgresql.org/docs/current/static/index.html.

Some sections contain sample commands that use the following default values:

- Default database name: activevos
- Default database user name: bpeluser
- Default database password: bpel

Backing up the PostgreSQL database on Linux

Use the script `db_backup.sh` to back up the PostgreSQL database.

To back up the PostgreSQL database, perform the following steps:

1. Go to `<Secure Agent installation directory>/apps/process-engine/data/db/util`.
2. Run the following command:
   `db_backup.sh <dbusername> <dbpassword> <path to backup file along with name of backup file>.dump`.

For example, the Secure Agent creates the backup file `backupfile1.dump` in the location `/home/data/myfolder/` if you run the following command:

```
db_backup.sh bpeluser bpel "/home/data/myfolder/backupfile1.dump",
```

Restoring the PostgreSQL database on Linux

Use the script `db_restore.sh` to restore the PostgreSQL database from a backup file.

To restore the PostgreSQL database file from a backup file, perform the following steps:

1. Go to `<Secure Agent installation directory>/apps/process-engine/data/db/util`.
2. Run the following command:
   `db_restore.sh <dbusername> <dbpassword> <path to dump file>`.

For example, you use the file `backupfile1.dump` to restore the PostgreSQL database if you run the following command:

```
db_restore.sh bpeluser bpel "/home/data/myfolder/backupfile1.dump",
```

Resetting the PostgreSQL database on Linux

You first shut down the PostgreSQL database and then use the script `db_reset.sh` to reset it.

To reset the PostgreSQL database to its original state, perform the following steps:

1. Go to `<Secure Agent installation directory>/apps/process-engine/data/db/util`
2. To shut down the server, run the following script:
   `server_stop.sh`
3. To reset the PostgreSQL database, run the following script:
   `db_reset.sh`

Starting the PostgreSQL server on Linux

To start the PostgreSQL Server, perform the following steps:

1. Go to `<Secure Agent installation directory>/apps/process-engine/data/db/util`.
2. Run the following script:
   `server_start.sh`.

Stopping the PostgreSQL server on Linux

To stop the PostgreSQL server, perform the following steps:

1. Go to `<Secure Agent installation directory>/apps/process-engine/data/db/util`.
2. Run the following script:
   `server_stop.sh`.

Getting the PostgreSQL server status on Linux

To get the status of the PostgreSQL server on Linux, perform the following steps:

1. Go to `<Secure Agent installation directory>/apps/process-engine/data/db/util`.
2. Run the following script:
   `server_status.sh`.

Vacuuming the PostgreSQL database on Linux

Vacuum the PostrgeSQL database to delete obsolete tuples and gain space. You use the script `db_maintenance.sh` to vacuum the PostgreSQL database.

By default, the PostgreSQL database auto vacuums. If you want to manually vacuum the database, perform the following steps:

1. Go to `<Secure Agent installation directory>/apps/process-engine/data/db/util`.
2. To vacuum the entire database, run the following command:
   `db_maintenance <dbusername> <dbpassword> vacuum`
3. To vacuum a single table, run the following command:
   `db_maintenance.sh <dbusername> <dbpassword> vacuum <tablename>`

For example, you vacuum the 'aeprocesslogdata' table if you run the following command:

```
db_maintenance.sh bpeluser bpel vacuum aeprocesslogdata
```

Reindexing the PostgreSQL database on Linux

Use the reindexing option to clean the index and free up space after you vacuum data on PostgreSQL. You use the script `db_maintenance.sh` to reindex the PostgreSQL database.

To reindex the PostgreSQL database, perform the following steps:

1. Go to `<Secure Agent installation directory>/apps/process-engine/data/db/util`.
2. To reindex the entire database, run the following command:
   `db_maintenance <dbusername> <dbpassword> reindex`
3. To reindex a single table, run the following command:
   `db_maintenance.sh <dbusername> <dbpassword> reindex <tablename>`

For example, you reindex the 'aeprocesslogdata' table if you run the following command:

```
db_maintenance.sh bpeluser bpel reindex aeprocesslogdata
```

Resetting transaction logs on Linux

If the PostgreSQL server does not start because of corruption to the control information, use the command `pg_resetxlog` to reset the control information.

To reset the control information of the PostgreSQL database, perform the following steps:

1. Go to `<Secure Agent installation directory>/apps/process-engine/data/db/postgresql-linux-x64-binaries/pgsql/bin .`
2. Run the following command:
   `pg_resetxlog -D <path to postgreSQL data directory>`

For example, you reset the transactions logs in the `Data` directory, if you run the following command:

```
pg_resetxlog -D "home/apps/process engine/data/PostGreSql/Data"
```

Configuring Secure Agent service properties

To configure Secure Agent service properties, open the **Runtime Environments** page and edit the Secure Agent. You can change or reset service property values and add and remove custom properties for a service. You can also change the Secure Agent name.

Note: Custom properties are specific to connectors. For more details about custom properties, see the help for the appropriate connector.

1. On the **Runtime Environments** page, click the name of the Secure Agent.

   Note: You might have to expand the Secure Agent group to see the list of Secure Agents within the group.
2. Click the **Details** tab.
3. In the upper right corner, click **Edit**.
4. To change the Secure Agent name, enter a new name in the **Agent Name** field.
5. To edit a service property, perform the following steps:
   a. In the **System Configuration Details** area, select a service.
   b. Select the configuration property type.
   c. In the row that contains the property that you want to change, click the **Edit Agent Configuration** icon and enter the new property value.
   d. To reset the property to the system default value, click the **Reset Agent Configuration to system default** icon.
6. To add a custom property for a service, perform the following steps:
   a. Scroll down to the **Custom Configuration Details** area.

      The following image shows the **Custom Configuration Details** area:

      

   b. Select the service that you want to configure.
   c. Select a configuration property type.
   d. If the configuration property type has a sub-type, select the appropriate sub-type.

      For example, to determine the logging level, choose INFO or DEBUG as the sub-type.
   e. Enter the property name and value.
   f. Click the **Add** icon.
7. To remove a custom property, click the **Remove** icon next to the custom property.
8. To reset all configuration properties to the default settings, click **Reset All**.
9. Click **Save**.

CHAPTER 14

Secure Agent installation

You can install a Secure Agent on Windows or Linux. You can also uninstall a Secure Agent if you no longer want to run the Secure Agent on the machine or if you want to reinstall the Secure Agent.

Installation and uninstallation instructions vary based on your operating system.

If you use the Secure Agent to run mapping tasks that are based on an elastic mapping, the Secure Agent must be installed on a Linux virtual machine on your cloud platform or you can use a serverless runtime environment.

Secure Agent installation on Windows

On Windows, the Secure Agent runs as a Windows service. When you install the Secure Agent, you also install the Informatica Cloud Secure Agent Manager.

By default, the Secure Agent starts when you start Windows. Use the Secure Agent Manager to stop and restart the Secure Agent. You can also use the Secure Agent Manager to check the Secure Agent status and configure proxy information.

You can launch the Secure Agent Manager from the Start menu or desktop icon. When you close the Secure Agent Manager, it minimizes to the Windows taskbar notification area for quick access.

When you install a Secure Agent, you perform the following tasks:

1. Verify that the machine where you install the Secure Agent meets the minimum requirements.
2. Download the Secure Agent installer files.
3. Install and register the Secure Agent.

Secure Agent requirements on Windows

You can install the Secure Agent on any machine that has internet connectivity and can access Informatica Intelligent Cloud Services.

Verify the following requirements before you install the Secure Agent on Windows:

- Verify that the machine on which you install the Secure Agent uses a supported operating system. For the list of supported operating systems for the Secure Agent, see the Product Availability Matrix (PAM) for Informatica Intelligent Cloud Services on the Product Availability Matrices page on Informatica Network.
- Verify that the machine where you install the Secure Agent has at least 5 GB of free disk space.
- Verify that the account you use to install the Secure Agent has access to all remote directories that contain flat source or target files.

- Verify that no other Secure Agent is installed on the machine. If another Secure Agent is installed on the machine, you must uninstall it first.

For more information about Secure Agent requirements, contact Informatica Global Customer Support.

Configure the firewall

If your organization uses a protective firewall, include the Informatica Intelligent Cloud Services domain name or IP address ranges in the list of approved domain names or IP addresses. You should also enable the port that the Secure Agent uses. This ensures that the Secure Agent can perform all necessary tasks through the firewall.

The Secure Agent uses port 443 (HTTPS) to connect to the internet. Configure your firewall to allow traffic to pass over port 443.

The whitelists of domains and IP addresses can vary according to your data center, which is also called a POD (Point of Deployment). You can identify your POD through the URL that appears when you open any service in Informatica Intelligent Cloud Services. The first few characters of the URL string identify the POD. For example, if the URL starts with `usw3.dm-us.informaticacloud.com`, your POD is USW3.

You can find the whitelists of Informatica Intelligent Cloud Services domains and IP addresses for different PODs in this Knowledge Base article on Informatica Network or by clicking the link at the top of the **Runtime Environments** page in Administrator.

Secure Agent permissions on Windows

A Secure Agent requires certain permissions to transfer data between sources and targets.

When you install a Secure Agent on Windows, the Secure Agent must be part of the local Administrators group.

Configuring Windows settings

Before you use the Secure Agent on Windows, configure proxy settings and a Windows Secure Agent service login.

You can configure proxy settings in Secure Agent Manager. Configure a login for the Windows Secure Agent service on Windows.

Note: If you use the Secure Agent for Informatica Cloud Data Wizard, you do not need to configure proxy settings or a Windows service login for the Secure Agent.

Downloading and installing the Secure Agent on Windows

To install the Secure Agent on a Windows machine, you must download and run the Secure Agent installation program and then register the agent.

Secure Agent registration requires an install token. To get the install token, copy the token when you download the agent or use the **Generate Install Token** option in Administrator. The token expires after 24 hours.

Before you download and install the Secure Agent, verify that no other Secure Agent is installed on the machine. If there is, you must uninstall it.

1. Open Administrator and select **Runtime Environments**.
2. On the **Runtime Environments** page, click **Download Secure Agent**.
3. Select the Windows 64-bit operating system platform, copy the install token, and then click **Download**.

The installation program is downloaded to your machine. The name of the installation program is `agent64_install_ng_ext.exe`.

4. Run the installation program:
   a. Specify the Secure Agent installation directory, and click **Next**.
   b. Click **Install** to install the agent.

   The Secure Agent Manager opens and prompts you to register the agent as shown in the following image:

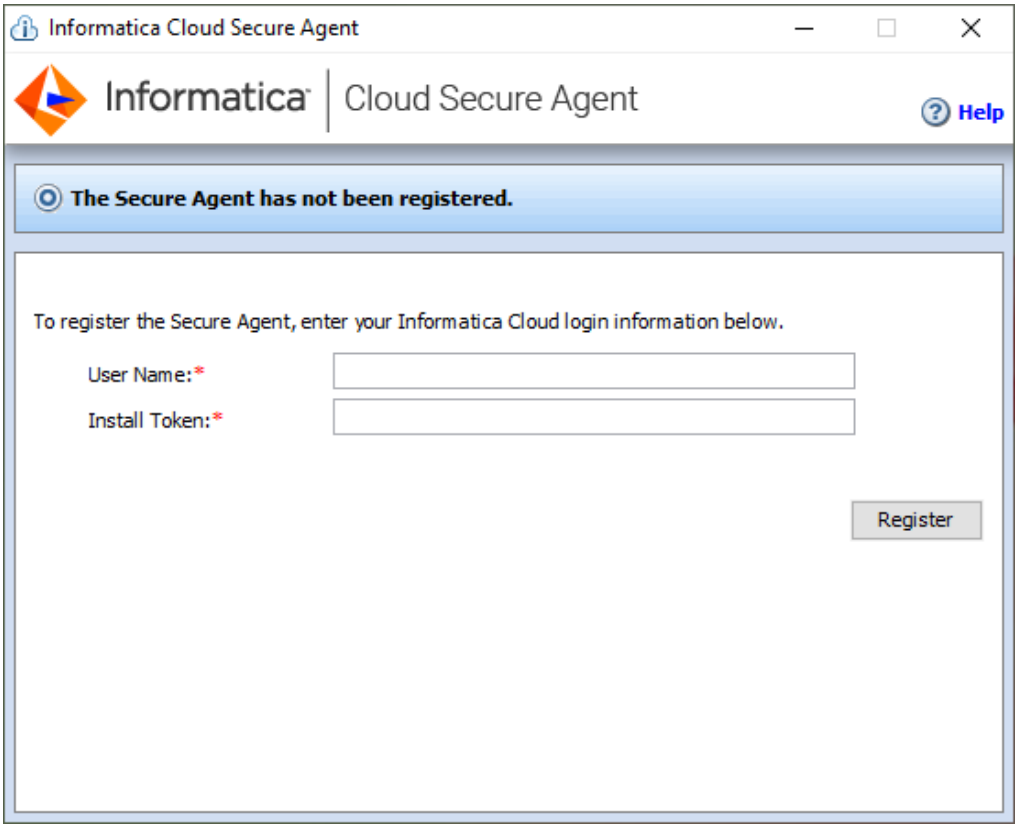

5. If you did not copy the install token when you downloaded the agent, click **Generate Install Token** on the **Runtime Environments** page in Administrator, and copy the token.
6. In the Secure Agent Manager, enter the following information, and then click **Register**:

Option	Description
User Name	User name that you use to access Informatica Intelligent Cloud Services.
Install Token	Token that you copied.

   The Secure Agent Manager displays the status of the Secure Agent. It takes a minute for all of the services to start.

7. If your organization uses an outgoing proxy server to connect to the internet, enter the proxy server information.
8. Close the Secure Agent Manager.

   The Secure Agent Manager minimizes to the taskbar and continues to run as a service until stopped.

Configure the proxy settings on Windows

If your organization uses an outgoing proxy server to connect to the internet, the Secure Agent connects to Informatica Intelligent Cloud Services through the proxy server. The Secure Agent installer configures the proxy server settings for the Secure Agent based on settings configured in the browser. You can change the proxy server settings through the Secure Agent Manager.

Contact your network administrator for the correct proxy settings.

1. In the Secure Agent Manager, click **Proxy**.
2. To enter proxy server settings, click **Use a Proxy Server**.
3. Enter the following information:

Field	Description
Proxy Host	Required. Host name of the outgoing proxy server that the Secure Agent uses.
Proxy Port	Required. Port number of the outgoing proxy server.
User Name	User name to connect to the outgoing proxy server.
Password	Password to connect to the outgoing proxy server.

4. Click **OK**.

   The Secure Agent Manager restarts the Secure Agent to apply the settings.

Configure a login for a Windows Secure Agent Service

On Windows, configure a network login for the Secure Agent service. The Secure Agent can access the network with the privileges and permissions associated with the login.

Configure a login for the machine on which the Secure Agent is installed to allow the Secure Agent to access directories to configure and run tasks. When you configure connections, configure tasks, and run tasks that use Flat File or FTP/SFTP connection types, the Secure Agent might require read and write permissions on the related directories.

For example, to browse to a directory to configure a Flat File or FTP/SFTP connection, the Secure Agent login might require permission to access the directory. Without a Secure Agent login with the appropriate permissions, Informatica Intelligent Cloud Services cannot display the directory in the **Browse for Directory** dialog box.

1. Go to the **Services** window from the Windows Administrative tools.
2. In the **Services** window, right-click the Informatica Cloud Secure Agent service and choose **Properties**.
3. In the **Properties** dialog box, click the **Log On** tab.
4. To configure a login, select **This Account**.
5. Enter an account and password.

   Use an account with the required privileges and permissions for the network security defined for the domain. By default, the account format is <domain name>\<user name>.
6. Click **OK**.
7. In the **Services** window, restart the Secure Agent service for the changes to take effect.

Uninstalling the Secure Agent on Windows

You can uninstall the Secure Agent. You might uninstall the Secure Agent if you no longer want to run the Secure Agent on the machine or if you want to reinstall the Secure Agent.

Before you uninstall the Secure Agent, verify that no connection or task is configured to use it.

1. Click **Start** > **All Programs** > **Informatica Cloud Secure Agent** > **Uninstall Informatica Cloud Secure Agent**.

   The Secure Agent uninstaller launches.
2. Click **Uninstall**.
3. When the uninstall completes, click **Done**.
4. Delete any remaining files in the installation directory.

   After you uninstall the Secure Agent, delete all files and directories associated with the Secure Agent installation.

Secure Agent installation on Linux

On Linux, the Secure Agent runs as a process. You can use a shell command line to install, register, start, stop, and uninstall the Secure Agent.

You can also use the shell command line to check the Secure Agent status.

When you install a Secure Agent, you perform the following tasks:

1. Verify that the machine where you install the Secure Agent meets the minimum requirements.
2. Download the Secure Agent installer files.
3. Install and register the Secure Agent.

Secure Agent requirements on Linux

You can install the Secure Agent on any machine that has internet connectivity and can access Informatica Intelligent Cloud Services. Before you install the Secure Agent on Linux, verify the system requirements.

Verify the following requirements before you install the Secure Agent on Linux:

- Verify that the machine on which you install the Secure Agent uses a supported operating system. For the list of supported operating systems for the Secure Agent, see the Product Availability Matrix (PAM) for Informatica Intelligent Cloud Services on the Product Availability Matrices page on Informatica Network.
- Verify that the machine where you install the Secure Agent has at least 5 GB of free disk space.
- The account that you use to install the Secure Agent must have access to all remote directories that contain flat source or target files.
- If you use PowerCenter, install the Secure Agent using a different user account than the account you used to install PowerCenter.

  Informatica Intelligent Cloud Services and PowerCenter use some common environment variables. If the environment variables are not set correctly for Informatica Intelligent Cloud Services, your jobs might fail at run time.

For more information about Secure Agent requirements, contact Informatica Global Customer Support.

Configure the firewall

If your organization uses a protective firewall, include the Informatica Intelligent Cloud Services domain name or IP address ranges in the list of approved domain names or IP addresses. You should also enable the port that the Secure Agent uses. This ensures that the Secure Agent can perform all necessary tasks through the firewall.

The Secure Agent uses port 443 (HTTPS) to connect to the internet. Configure your firewall to allow traffic to pass over port 443.

The whitelists of domains and IP addresses can vary according to your data center, which is also called a POD (Point of Deployment). You can identify your POD through the URL that appears when you open any service in Informatica Intelligent Cloud Services. The first few characters of the URL string identify the POD. For example, if the URL starts with `usw3.dm-us.informaticacloud.com`, your POD is USW3.

You can find the whitelists of Informatica Intelligent Cloud Services domains and IP addresses for different PODs in this Knowledge Base article on Informatica Network or by clicking the link at the top of the **Runtime Environments** page in Administrator.

Secure Agent permissions on Linux

A Secure Agent requires certain permissions to transfer data between sources and targets.

When you install a Secure Agent on Linux, the Secure Agent must have read/write/execute permissions for the installation directory.

Downloading and installing the Secure Agent on Linux

To install the Secure Agent on a Linux machine, you must download and run the Secure Agent installation program and then register the agent.

Secure Agent registration requires an install token. To get the install token, copy the token when you download the agent or use the **Generate Install Token** option in Administrator. The token expires after 24 hours.

Before you download and install the Secure Agent, verify that no other Secure Agent is installed on the machine using the same Linux user account. If there is, you must uninstall it.

1. Open Administrator and select **Runtime Environments**.
2. On the **Runtime Environments** page, click **Download Secure Agent**.
3. Select the Linux 64-bit operating system platform, copy the install token, and then click **Download**.

   The installation program is downloaded to your machine. The name of the installation program is `agent64_install_ng_ext.bin`.
4. Save the installation program to a directory on the machine where you want to run the Secure Agent.

   Note: If the file path contains spaces, the installation might fail.
5. From a shell command line, navigate to the directory where you downloaded the installation program and enter the following command:

   ```
   ./agent64_install_ng_ext.bin -i console
   ```
6. When the installer completes, navigate to the following directory:

   ```
   <Secure Agent installation directory>/apps/agentcore
   ```
7. To start the Secure Agent, enter the following command:

   ```
   ./infaagent startup
   ```

The Secure Agent Manager starts. You must register the agent using the user name that you use to access Informatica Intelligent Cloud Services. You must also supply the install token.

8. If you did not copy the install token when you downloaded the agent, click **Generate Install Token** on the **Runtime Environments** page in Administrator, and copy the token.
9. In the `<Secure Agent installation directory>/apps/agentcore` directory, enter the following command using your Informatica Intelligent Cloud Services user name and the token that you copied:

   ```
   ./consoleAgentManager.sh configureToken <user name> <install token>
   ```

   You can check the registration status of a Secure Agent using the following command:

   ```
   ./consoleAgentManager.sh isConfigured
   ```

Configure the proxy settings on Linux

If your organization uses an outgoing proxy server to connect to the internet, the Secure Agent connects to Informatica Intelligent Cloud Services through the proxy server.

The Secure Agent installer configures the proxy server settings for the Secure Agent based on settings configured in the browser. You can update the proxy server settings defined for the Secure Agent from the command line.

To configure the proxy server settings for the Secure Agent on a Linux machine, use a shell command that updates the `proxy.ini` file. Contact the network administrator to determine the proxy settings.

1. Navigate to the following directory:

   ```
   <Secure Agent installation directory>/apps/agentcore
   ```

2. To update the `proxy.ini` file, enter the following command:

   ```
   ./consoleAgentManager.sh configureProxy <proxy host> <proxy port> <proxy user name>
   <proxy password>
   ```

3. Restart the Secure Agent.

Uninstalling the Secure Agent on Linux

You can uninstall the Secure Agent. You might uninstall the Secure Agent if you no longer want to run the Secure Agent on the machine or if you want to reinstall the Secure Agent.

Before you uninstall the Secure Agent, verify that no connection or task is configured to use it.

1. From the command line, navigate to the following directory:

   ```
   <Secure Agent installation directory>/apps/agentcore
   ```

2. Stop the Secure Agent Linux process by entering the following command:

   ```
   ./infaagent shutdown
   ```

3. To uninstall the Secure Agent, run rm -rf on the directory where you installed the Secure Agent to remove Secure Agent files.

CHAPTER 15

Schedules

You can create schedules to run tasks or taskflows at specified times or at regular intervals. You can also define a blackout period during which scheduled tasks or jobs do not run.

Create schedules and configure blackout periods on the **Schedules** page in Administrator. After you create a schedule, you can associate it with tasks and taskflows in another service such as Data Integration.

When you create a schedule, you specify the date and time. You can configure a schedule to run associated assets throughout the day between 12:00 a.m. and 11:55 p.m. Informatica Intelligent Cloud Services might add a small schedule offset to the start time, end time, and all other time configurations. As a result, scheduled tasks and taskflows might start later than expected. For example, you configure a schedule to run hourly until noon, and the schedule offset for your organization is 10 seconds. Informatica Intelligent Cloud Services extends the end time for the schedule to 12:00:10 p.m., and the last hourly task or taskflow starts at 12:00:10 p.m. To see the schedule offset for your organization, check the **Schedule Offset** organization property for the Data Integration service.

You can perform the following tasks with schedules:

Associate a schedule with a task or taskflow

To associate a schedule with a task or taskflow, edit the task or taskflow. For example, to associate a schedule with a mapping task, edit the mapping task in Data Integration, and select the schedule on the **Schedules** page.

When you copy a task or taskflow that includes a schedule, the schedule is not associated with the new asset. To associate a schedule with the new asset, edit the asset.

Monitor scheduled tasks

You can monitor scheduled tasks from the **All Jobs** page in Monitor. Scheduled tasks do not appear on the **My Jobs** page.

Export a schedule

You can export a schedule from your organization and import it into another organization. Export a schedule on the **Schedules** page. If the schedule is associated with a task or taskflow, the task or taskflow is not included in the export file.

Delete a schedule

Delete a schedule on the **Schedules** page.

Note: You cannot delete a schedule that is used in a task or taskflow. Remove the schedule from all tasks and taskflows before you delete the schedule.

Configuring a blackout period

A blackout period prevents all scheduled tasks and taskflows in the organization from running during a specified period of time. You can configure one blackout period for an organization.

To configure a blackout period, in Administrator, select **Schedules**, and then click **Blackout Period**. The blackout period is displayed on the **Schedules** page.

Repeat frequency

The repeat frequency determines how often tasks run. The following table describes the repeat frequency options:

Option	Description
Does not repeat	Tasks run as scheduled and do not repeat.
Every N minutes	Tasks run on an interval based on a specified number of minutes. You can configure the following options: - Repeat frequency. Select a frequency in minutes. Options are 5, 10, 15, 20, 30, 45. - Days. Days of the week when you want tasks to run. You can select one or more days of the week. - Time range. Hours of the day when you want tasks to start. Select All Day or configure a time range. You can configure a time range between 00:00-23:55. - Repeat option. The range of days when you want tasks to run. You can select Repeat Indefinitely or configure an end date and time.
Hourly	Tasks run on an hourly interval based on the start time of the schedule. You can configure the following options: - Repeat frequency. Select a frequency in hours. Options are 1, 2, 3, 4, 6, 8, 12. - Days. Days of the week when you want tasks to run. You can select one or more days of the week. - Time range. Hours of the day when you want tasks to start. Select All Day or configure a time range. You can configure a time range between 00:00-23:55. - Repeat option. The range of days when you want tasks to run. You can select Repeat Indefinitely or configure an end date and time.
Daily	Tasks run daily at the start time configured for the schedule. You can configure the following options: - Repeat frequency. The frequency at which you want tasks to run. Select Every Day or Every Weekday. - Repeat option. The range of days when you want tasks to run. You can select Repeat Indefinitely or configure an end date and time.
Weekly	Tasks run on a weekly interval based on the start time of the schedule. You can configure the following options: - Days. Days of the week when you want tasks to run. You can select one or more days of the week. - Repeat option. The range of days when you want tasks to run. You can select Repeat Indefinitely or configure an end date and time. If you do not specify a day, the schedule runs regularly on the same day of the week as the start date.

Option	Description
Biweekly	Tasks run every two weeks based on the start time of the schedule. You can configure the following options: - Days. Days of the week when you want tasks to run. You can select one or more days of the week. You must select at least one day. - Repeat option. The range of days when you want tasks to run. You can select Repeat Indefinitely or configure an end date and time. If you configure a biweekly schedule to start at 5 p.m. on a Tuesday and run tasks every two weeks on Mondays, the schedule begins running tasks on the following Monday.
Monthly	Tasks run on a monthly interval based on the start time of the schedule. You can configure the following options: - Day. Day of the month when you want tasks to run. You can configure one of the following options: - Select the exact date of the month, between 1-28. If you want the task to run on days later in the month, use the <n> <day of the week> option. - Select the <n> <day of the week>. Options for <n> include First, Second, Third, Fourth, and Last. Options for <day of the week> includes Day, and Sunday-Saturday. Tip: With the Day option, you can configure tasks to run on the First Day or the Last Day of the month. - Repeat option. The range of days when you want tasks to run. You can select Repeat Indefinitely or configure an end date and time.

Time zones and schedules

Informatica Intelligent Cloud Services stores time in Coordinated Universal Time (UTC). When you log in, Informatica Intelligent Cloud Services converts the time and displays it in the time zone associated with your user profile.

When you create a schedule, you select the time zone for the scheduler to use. You can select a time zone that is different from your time zone or your organization time zone.

Daylight Savings Time changes and schedules

Informatica Intelligent Cloud Services applies Daylight Savings Time changes applies Daylight Savings Time changes to all tasks except biweekly tasks.

When Daylight Savings time goes into effect, tasks scheduled to run between 2:00 a.m. and 2:59 a.m., do not run the day that the time changes from 2:00 a.m. to 3:00 a.m. If a task is scheduled to run biweekly at 2 a.m., it will run at 3 a.m. the day of the time change and at 2 a.m. for the next run.

Daylight Savings Time does not trigger additional runs for tasks that are scheduled to run between 1:00 a.m. - 1:59 a.m. when Standard Time begins. For example, a task is scheduled to run every day at 1:30 a.m. When the time changes from 2 a.m. to 1 a.m., the task does not run again at 1:30 a.m.

Tip: To ensure that Informatica Intelligent Cloud Services does not skip any scheduled runs near the 2 a.m. time change, do not schedule jobs to run between 12:59 a.m. and 3:01 a.m.

Configuring a schedule

Configure a schedule on the **Schedules** page. For mapping and synchronization tasks, you can also create a new schedule when you configure the task. You can configure a schedule to run once or at a specific interval and to run indefinitely or until a specified end time.

1. In Administrator, select **Schedules**.
2. To create a schedule, click **New Schedule**.

   To edit a schedule, click the edit icon in the row that contains the schedule.
3. Configure the following properties:

Property	Description
Schedule Name	Name of the schedule. Each schedule name must be unique within the organization. Schedule names can contain alphanumeric characters, spaces, and the following special characters: _ . + - Maximum length is 100 characters. Names are not case sensitive.
Description	Description of the schedule. Maximum length is 255 characters.
Starts	Date and time when the schedule starts. The date format is MM/DD/YYYY. Time appears in the 24-hour format. Click the calendar button to select the start date. The start date and time can affect the repeat frequency for tasks and taskflow jobs that repeat at regular intervals. For example, if the start date is November 10 and the repeat frequency is monthly, the schedule runs associated assets on the tenth day of each month. If the start time is 3:10 and the repeat frequency is hourly, the assets run every hour at 10 minutes past the hour. Default is the current date, current time, and time zone of the user who creates the schedule.
Time Zone	Select the time zone for the schedule to use. The time zone can differ from the organization time zone or user time zone.
Repeats	Repeat frequency for the schedule. Select one of the following options: - Does Not Repeat - Every N Minutes - Hourly - Daily - Weekly - Biweekly - Monthly Default is Does Not Repeat.

4. Click **Save**.

Exporting schedules

You can export schedules from your organization and import them into other organizations. Assets that are associated with the schedules are not included in the export file. Export schedules on the **Schedules** page.

1. In Administrator, select **Schedules**.
2. Click **Export**.
3. In the **Export Schedules** dialog box, select the schedules that you want to export.
4. Optionally, update the export job name.

   By default, the job name is `SchedulesExport_<date>`.
5. Click **Export**.

   Administrator creates an export job to export the schedule.
6. To check the status of the export job and download the export file, open the **Import/Export Logs** page in Monitor, and click the **Export** tab.

   You can download the export file in the row that contains the export job or on the job details page.

You can import schedules on the **Explore** page in another service such as Data Integration. For information about importing assets, see the help for that service.

After you import schedules, you can associate them with assets in the target organization.

CHAPTER 16

Bundle management

A bundle is a set of related mappings, mapping tasks, mapplets, and Visio templates that Data Integration users can use in data integration projects. Data Integration users design, create, and publish bundles. Administrators manage bundles.

If you are the administrator for an organization, you can perform the following actions to manage bundles:

Install a bundle.

You can install a public, private, or unlisted bundle that the bundle designer has configured to be used as a reference. When you install a bundle, the bundle is installed into the Add-On Bundles project in Data Integration. Users in your organization can use the assets in the bundle, but they cannot edit the assets.

Copy a bundle.

You can copy a public, private, or unlisted bundle that the bundle designer has configured for copying. When you copy a bundle, you select the Data Integration folder where you want to copy the bundle contents. You can copy a bundle multiple times and save the contents into a different project or folder each time that you copy it. After you copy a bundle, users in your organization can edit the assets.

Upgrade a bundle.

If you installed a bundle and a newer version of the bundle is available, you can upgrade the bundle to get the latest version.

Uninstall a bundle.

If your organization no longer needs an installed bundle, you can uninstall it.

To view the bundles that are installed or are available to your organization, in Administrator, select **Add-On Bundles**. The **Add-on Bundles** page displays information about installed bundles, copied bundles, and bundles that are available for installation or copying.

For information about bundle types, creating bundles, or publishing bundles, see *Mappings* in the Data Integration service help.

Installing a bundle

You can install a public, private, or unlisted bundle that the bundle designer has configured to be used as a reference. Install a bundle on the Available Bundles tab of the **Add-On Bundles** page.

Before you install an unlisted bundle, get the bundle access code. To get the access code for a bundle that was created in your organization, open the **Bundles** page in Data Integration, click the bundle name, and then

click **Copy Access Code**. To get the bundle access code for a bundle that was created outside of your organization, contact the bundle publisher.

1. In Administrator, select **Add-On Bundles**.
2. Click **Available Bundles**.

   The Available Bundles tab lists the public and private bundles that are available for installation or copying.
3. If the bundle that you want to install is an unlisted bundle, enter the bundle access code in the **Find** field.
4. Click the bundle name to open the Bundle Details page.
5. Verify that the **Allow** field is set to **Reference** or to **Reference and Copy**.

   You cannot install a bundle that is configured for copying only.
6. Click **Install**.

In Data Integration, the bundle is added to the Add-On Bundles project, and the assets are ready for use. The bundle is also listed on the **Installed Bundles** tab of the **Add-On Bundles** page in Administrator.

Copying a bundle

You can copy a public, private, or unlisted bundle that the bundle designer has configured for copying. Copy a bundle on the Available Bundles tab of the **Add-On Bundles** page. Each time you copy a bundle, an event is logged on the Copied Bundles tab.

Before you copy an unlisted bundle, get the bundle access code. To get the access code for a bundle that was created in your organization, open the **Bundles** page in Data Integration, click the bundle name, and then click **Copy Access Code**. To get the bundle access code for a bundle that was created outside of your organization, contact the bundle publisher.

1. In Administrator, select **Add-On Bundles**.
2. Click **Available Bundles**.

   The Available Bundles tab lists the public and private bundles that are available for installation or copying.
3. If the bundle that you want to copy is an unlisted bundle, enter the bundle access code in the **Find** field.
4. Click the bundle name to open the Bundle Details page.
5. Verify that the **Allow** field is set to **Copy** or to **Reference and Copy**.

   You cannot copy a bundle that is configured to be used as a reference only.
6. Click **Copy Bundle Content to...**.
7. In the **Browse** dialog box, select the Data Integration project or folder into which you want to copy the bundle contents.
8. Click **Select**.

   The assets in the bundle are copied to the selected project or folder.

Upgrading a bundle

You can upgrade an installed bundle when an updated version becomes available. You can check the bundle status on the Installed Bundles tab of the **Add-On Bundles** page.

1. In Administrator, select **Add-On Bundles**.
2. Click **Installed Bundles**.

   The Bundle Status column indicates whether the bundle is up-to-date or whether an upgrade is available.
3. Click the bundle name to open the Bundle Details page.
4. Click **Upgrade**.

Uninstalling a bundle

Uninstall a bundle if users in your organization no longer need it. Uninstall the bundle on the Installed Bundles tab of the **Add-On Bundles** page.

Note: Uninstalling a bundle removes all of the bundle assets from the organization. If you want to keep tasks that use assets in the bundle, remove the assets from the task before you uninstall the bundle.

1. In Administrator, select **Add-On Bundles**.
2. Click **Installed Bundles**.
3. Click the bundle name to open the Bundle Details page.
4. Click **Uninstall**.

After you uninstall the bundle, it is listed on the Available Bundles tab.

CHAPTER 17

Event monitoring

You can monitor events for the assets, licenses, users, and Secure Agents in your organization through the asset and security logs. To view the logs, you must be assigned a role that has the Audit Log - View privilege.

You can monitor events through the following logs:

Asset log

Displays the following information:

- Events for assets such as when an asset was created, updated, copied, or deleted and the name of the user who modified the asset.
- Authentication events for users such as when a user in the organization logged in to Informatica Intelligent Cloud Services.
- Events related to licenses such as when a license was added, removed, or changed.

To open the asset log, in Administrator, select **Logs**, and then select **Asset Logs** at the top of the page.

Security log

Displays events for Secure Agents and organizations such as when each agent was created or updated, when organization information was updated, and the name of the user who modified the agent or organization.

To open the security log, in Administrator, select **Logs**, and then select **Security Logs** at the top of the page.

The following image shows the asset log:

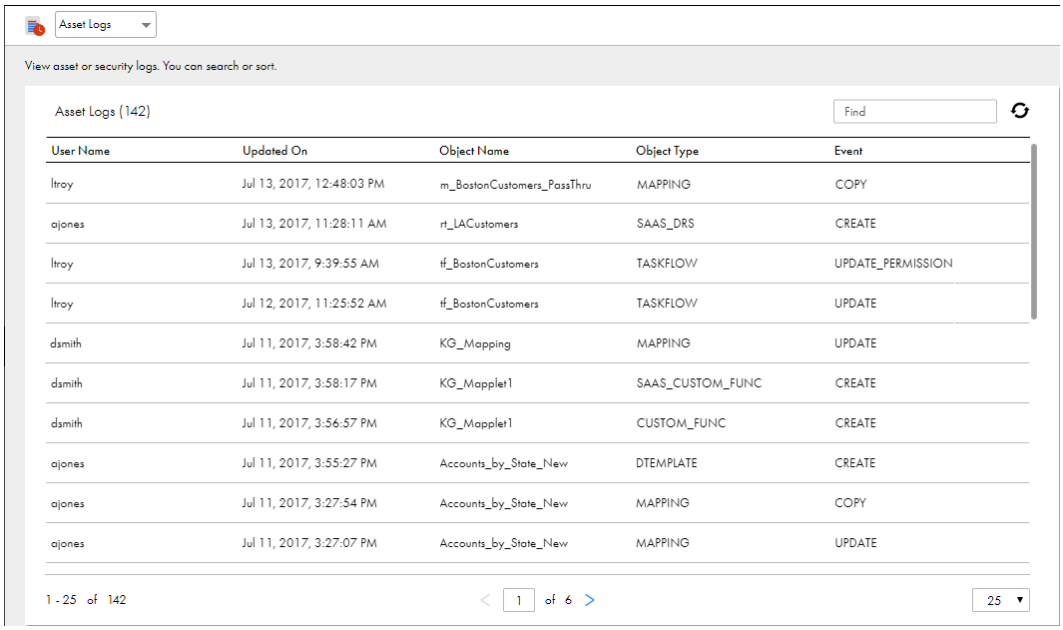

By default, the logs display events for the past 90 days. To change the length of time that events appear in the logs, contact Informatica Global Customer Support.

You can customize the properties that are displayed in the logs in the following ways:

- To hide a column, right-click the column heading area and uncheck the column that you want to hide.
- To sort the log events, click the column heading for the property that you want to sort by. To reverse the sort order, click the column heading again.
- To search the logs for specific events, enter the search string in the **Find** field. You can search for an object name or event type.

CHAPTER 18

File transfer

You can use file transfer protocols such as HTTPS, AS2, and SFTP to exchange files with remote partners.

To exchange files, you can use B2B Gateway or you can use the Data Integration REST API sendfiles resource.

To exchange files with a remote partner, configure your organization's file servers associated with the File Integration Service to securely communicate with the partner's servers. The File Integration Service is a Secure Agent service that runs advanced file transfer protocols.

You can configure the following types of file servers:

AS2 server

To receive files from partners with AS2 file transfer, configure an AS2 server to receive files from remote AS2 servers.

To send AS2 files to a partner's server, configure a connection, and then send the files to the partner using the Informatica Intelligent Cloud Services REST API. For more information, see the help for the AS2 Connector in the Data Integration help.

For example, you want to exchange EDI messages with a partner's AS2 server. To receive files from the partner, you configure your file server to accept files from the partner's server. To send files to your partner's server, you configure an AS2 connection for the partner. Then you send a POST request using the sendfiles REST API resource to transfer the EDI messages to the partner's server.

HTTPS server

To exchange files with partners with HTTPS file transfer, configure an HTTPS server so that partners can connect to the server, upload files to the server, and download files from the server.

SFTP server

To exchange files with partners with SFTP file transfer, configure an SFTP server so that partners can connect to the server, and upload to and download files from the server.

Proxy server

You can install and configure one or more file integration proxy servers in the demilitarized zone (DMZ). The partners' servers then communicate with the proxy servers instead of communicating directly with the organization's file servers. Multiple file servers can use the same file integration proxy server.

You can install proxy servers on Windows and Linux operating systems.

For each remote partner that exchanges files with your organization, you create a file server user account. You define the protocol accessibility for the file server user, that is, AS2, HTTPS, SFTP, or various combinations of these servers. A home directory is created or assigned to each file server user. You can define network shared locations to the user's home directory, and define folder level and file level permissions for the user.

You can monitor file transfer jobs on the **File Transfer Logs** page in Monitor. For information about monitoring file transfer jobs, see the Monitor help.

To exchange files, you must have the following licenses:

- File Integration Service
- HTTPS Server if you want to exchange HTTPS files
- AS2 Server and AS2 Connector if you want to exchange AS2 files
- SFTP Server if you want to exchange SFTP files

File server configuration process

Configure file servers, file server users, and global settings to exchange files between remote partners and your Informatica Intelligent Cloud Services organization.

You can configure file servers for each Secure Agent that uses the File Integration Service. Configure file servers on the **File Servers** page in Administrator. The **File Servers** page lists all of the runtime environments and Secure Agents in your organization that can use the File Integration Service.

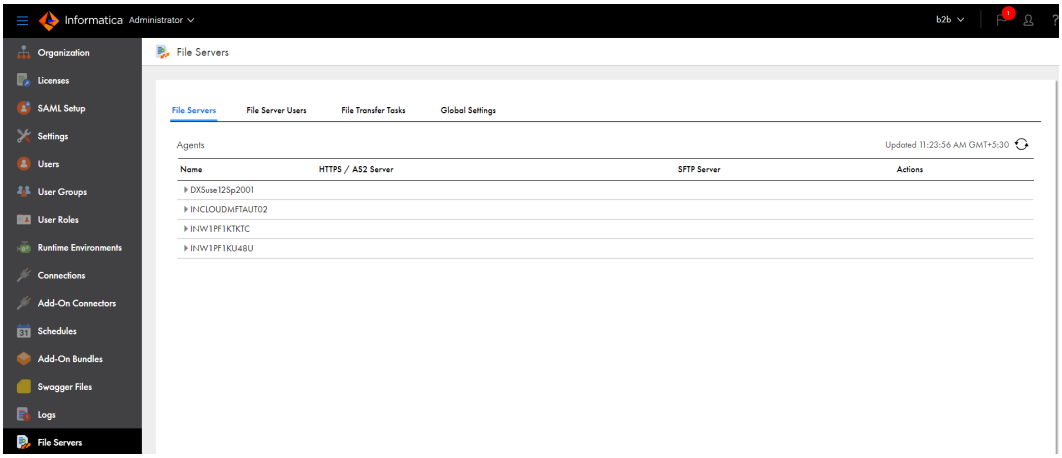

Enabling external partners to exchange files with your organization includes the following tasks:

- Configure file server properties. You can configure an AS2, HTTPS, and SFTP server.
- Optionally, install and configure one or more proxy servers as intermediaries between the partner's file servers and the organization's file servers.
- Configure server users for remote partners so that they can exchange files with your server.
- Specify default folders where the files are exchanged.

Before you begin

Before you configure file servers, ensure that you have the appropriate licenses and that you exchange public keys with the partner.

To ensure that your organization can exchange files with a remote server, complete the following tasks:

1. Send your public keys to the partner.
2. Receive the partner public keys.

3. Import the partner public keys to your trust store.
4. Configure the file server settings.

Ensure that the File Integration Service is running on the Secure Agent. For information about checking the status of Secure Agent services, see Chapter 13, “Secure Agent services” on page 114.

File servers

Configure file servers to exchange files with remote partners.

You can configure the following servers:

- AS2 server. Receives files from partners with AS2 file transfer.
- HTTPS server. Partners connect to the server to upload and download files.
- SFTP server. Partners connect to the server to upload and download files.
- Proxy server. An intermediary between the partner's file servers and the organization's file servers.

Configuring a file server

Configure properties for a file server to exchange files between the server and remote partners.

1. In Administrator, select **File Servers**.
2. On the **File Servers** tab, select the Secure Agent that runs the File Integration Service that you want to use to exchange files with the remote servers.
3. On the **File Server for agent** page, select the tab for the type of server to configure, HTTPS server, AS2 server, SFTP server, or proxy server.
4. Configure the file server properties, and then click **Save**.
   - For information about AS2 server properties, see “AS2 server configuration properties” on page 165.
   - For information about HTTPS server properties, see “HTTPS server configuration properties” on page 169.
   - For information about SFTP server properties, see “SFTP server configuration properties” on page 172.
   - For information about proxy server properties, see “Proxy server configuration properties” on page 174.

AS2 server configuration properties

For each runtime environment that uses the File Integration Service, you can configure an AS2 server to receive files from remote AS2 servers.

You configure AS2 server properties on the **AS2 Server** tab of the **File Server for agent** page.

Configure the following types of properties:

- General properties
- SSL properties
- Message security properties

- MDN properties
- Upload restriction properties

General properties

The following table describes general AS2 server properties:

Property	Description
Enable AS2 Server	Whether to enable the AS2 server. When not enabled, the AS2 server cannot receive files. Default is disabled.
AS2 Server Id	Name or ID used by the sender. Note the following rules for the ID: - The value is case sensitive. - The ID can contain up to 128 ASCII characters, special characters, and spaces.
Port	Port number for the AS2 server. Default is 15400.
Local Address	Local address of the AS2 sever.
Enable SSL	Whether to use SSL encryption in communications with remote AS2 servers. Default is disabled.

SSL properties

The following table describes the SSL properties:

Property	Description
SSL Protocol	Whether to use the SSL or TLS protocol. Select one of the following values: - TLS. A new version of SSL, Transport Layer Security will be used to secure the transmission. - SSL. A traditional Secure Socket Layer protocol is used to secure the transmission. Default is SSL.
Enabled SSL Protocols	Specify the permissible TLS and SSL versions separated with a comma. The supported versions are: - TLS: TLSv1.2, and TLSv1.3 - SSL: SSLv2Hello and SSLv3 When a value is not specified, all the versions for the selected protocol are enabled.
Client Authentication	Whether the client must have a certificate to authenticate with the server. Choose one of the following values: - None. The SSL connection runs without checking certificates and the user is authenticated with a password. If any of the information being transmitted requires a certificate, the connection fails. - Required. The SSL connection will not connect or authenticate a user unless a valid certificate is available. - Optional. The SSL connection looks for a valid certificate, but continues with password authentication if a certificate is not present.

Property	Description
Key Store Location	Location of the key store that stores the private key and associated certificates that the client uses to authenticate communications with the File Integration Service. Include path and file name.
Key Store Password	Password to access the key store.
Key Store Type	Type of the private key store. Use one of the following values: - JKS - PKCS12
Key Alias	Key alias or certificate for the private key used to sign the MDN.
Trust Store Location	The path to the trust store file that the File Integration Service uses for HTTPS communication.
Trust Store Password	Password to access the trust store.
Trust Store Type	Type of trust store. Use one of the following values: - JKS - PKCS12

Message security properties

The following table describes basic message security properties:

Property	Description
Encryption Required	Whether files received by the File Integration Service must be encrypted. Default is enabled.
Signature Required	Whether files from the remote AS2 server must contain a digital signature. If a signature is required, the File Integration Service rejects any messages without the signature. Default is enabled.
Authentication Required	Whether the user is required to authenticate. Default is disabled.
Decryption Certificate Alias	Key alias or certificate used to decrypt incoming messages. The alias references a certificate in the key store. All partners who send AS2 messages must have the public portion of this certificate.

MDN properties

The following table describes message receipt properties:

Property	Description
MDN Signature Certificate Alias	Alias that refers to the private key that the AS2 server uses to sign the message receipt. The private key is in the default private key store.
Asynchronous MDN Automatic Approval	Whether to send a return receipt automatically or manually.
Enabled Proxy for Async MDN	Determines if a proxy server is enabled for Asynchronous MDN. Default is disabled.
Proxy Type	Type of proxy server to use for the connection. Select one of the following types: - SOCKS. You can use SOCKS version 4 or 5. - HTTPS. - Informatica File Server proxy. Verify with your network administrator which proxy server type to use.
Host	Host name or IP address of the proxy server on your network.
Port	Port number of the proxy server on your network. If left blank, the default port for HTTP is 80 and the default port for SOCKS is 1080.
User	User name to use for login when connecting to the proxy server.
Password	Password for connecting to the proxy server. Required if your network uses the proxy server to create HTTP or HTTPS connections.

Upload restrictions properties

You can specify the types of files to allow or deny in an AS2 file upload. The following table describes the properties that control upload restrictions:

Property	Description
File Extension Filter Type	Whether to accept or deny the extensions in the File Extensions list. Use one of the following values: - Do Not Filter. Accept all file types. - Accept. Accept files with the extensions listed in the File Extensions property. - Deny. Do not allow files with the extensions listed in the File Extensions property.
File Extensions	List of file extensions. Add the file extensions that correspond to the File Extension Filter Type. For example, to accept .csv and .txt files, in the File Extension Filter Type property, select **Accept**, and then add csv and txt to the list of file extensions. To add an extension to the list, type the extension in the text box and click **Add**. To remove an extension from the list, highlight the extension and click **Delete**.
File Extension Case Sensitive	Whether to factor case when you filter using the file extensions list. When enabled, files with extensions that do not match the case used in the File Extensions list cannot be uploaded. For example, if the file extension list includes csv but not CSV, files with the extension of csv can be uploaded but files with the extension of CSV cannot be uploaded.

Property	Description
Allow Files with Extension	Whether to enable the file extension filter. When enabled, the file extension properties configured on this page determine which file types can be uploaded. Default is enabled.
Allow Files with No Extension	Whether to allow files that do not include the extension in the file name. Default is enabled.
Allow Files with No Name	Whether to allow files with no name. The Secure Agent saves files without a name using the following format: `as2data_<datetime>` where datetime is the current time stamp including milliseconds. Default is enabled.
File Name Suffix Timestamp (Optional)	Whether to append timestamp to the file name. When enabled, the timestamp is suffixed to the file name.
Max Upload Size	Maximum file size that the AS2 server can upload, in megabytes. Default is 5 MB.
When File Exists	Choose the action to be performed when a file that already exists in the folder is received again. Select one of the following options: - Rename: Rename the newly received file. - Append: Append the changes to the existing file. - Overwrite: Overwrite the existing file with the newly received file. - Error: Display an error if the file already exists.

HTTPS server configuration properties

For each runtime environment that uses the File Integration Service, you can configure an HTTPS server to exchange files with remote HTTPS servers.

You configure HTTPS server properties on the **HTTPS Server** tab of the **File Server for agent** page. You must have the HTTPS license to exchange files through HTTPS servers.

Configure the following types of properties:

- General
- SSL
- Upload restriction

General properties

The following table describes general HTTPS server properties:

Property	Description
Enable HTTPS Server	Whether to enable the HTTPS server. When not enabled, the HTTPS server cannot receive files. Default is disabled.
Port	Port number for the HTTPS server. Default is 15400.
Local Address	Local address of the HTTPS sever.
Enable SSL	Whether to use SSL encryption in communication with remote HTTPS servers. Default is disabled.

SSL properties

The following table describes the SSL properties:

Property	Description
SSL Protocol	Whether to use the SSL or TLS protocol. Select one of the following values: - TLS. A new version of SSL, Transport Layer Security will be used to secure the transmission. - SSL. A traditional Secure Socket Layer protocol is used to secure the transmission (default).
Enabled SSL Protocols	Specify the permissible TLS and SSL versions separated with a comma. The supported versions are: - TLS: TLSv1.1,TLSv1.2, and TLSv1.3 - SSL: SSLv2Hello and SSLv3 When a value is not specified, all the versions for the selected protocol are enabled.
Client Authentication	Whether the client must have a certificate to authenticate with the server. Choose one of the following values: - None. The SSL connection runs without checking certificates and the user is authenticated with a password. If any information being transmitted requires a certificate, the connection fails. - Required. The SSL connection will not connect or authenticate a user unless a valid certificate is available. - Optional. The SSL connection looks for a valid certificate, but continues with password authentication if a certificate is not present.
Key Store Location	Location of the key store that stores the private key and associated certificates. Client uses the key store file to authenticate communication with the File Integration Service. Include path and file name.
Key Store Password	Password to access the key store.

Property	Description
Key Store Type	Type of the private key store. Use one of the following values: - JKS - PKCS12
Key Alias	Key alias or certificate for the private key used to sign the MDN.
Trust Store Location	The path to the trust store file that the File Integration Service uses for HTTPS communication.
Trust Store Password	Password to access the trust store.
Trust Store Type	Type of trust store. Use one of the following values: - JKS - PKCS12

Upload restrictions properties

You can specify the types of files to allow or deny in an HTTPS file upload.

The following table describes the properties that control upload restrictions:

Property	Description
File Extension Filter Type	Whether to accept or deny the extensions in the File Extensions list. Use one of the following values: - Do Not Filter. Accept all file types. - Accept. Accept files with the extensions listed in the File Extensions property. - Deny. Do not allow files with the extensions listed in the File Extensions property.
File Extensions	List of file extensions. Add the file extensions that correspond to the File Extension Filter Type. For example, to accept .csv and .txt files, in the File Extension Filter Type property, select **Accept**, and then add csv and txt to the list of file extensions. To add an extension to the list, type the extension in the text box, and click **Add**. To remove an extension from the list, highlight the extension, and click **Delete**.
File Extension Case Sensitive	Whether to factor case when you filter using the file extensions list. When enabled, files with extensions that do not match the case used in the File Extensions list cannot be uploaded. For example, if the file extension list includes csv but not CSV, files with the extension of csv can be uploaded but files with the extension of CSV cannot be uploaded.
Allow Files with Extension	Whether to enable the file extension filter. When enabled, the file extension properties configured on this page determine which file types can be uploaded. Default is enabled.
Allow Files with No Extension	Whether to allow files that do not include the extension in the file name. Default is enabled.

Property	Description
Allow Files with No Name	Whether to allow files without a name. The Secure Agent saves files without a name using the following format: `as2data_<datetime>` where datetime is the current time stamp including milliseconds. Default is disabled.
Max Upload Size(MB)	The file size limit in megabytes for the HTTPS server upload. Default is 5 MB.

SFTP server configuration properties

For each runtime environment that uses the File Integration Service, you can configure an SFTP server to exchange files.

You configure SFTP server properties on the **SFTP Server** tab of the **File Server for agent** page. Configure the following types of properties:

- General properties
- Algorithms properties
- Host keys properties
- Upload restriction properties

General properties

The following table describes general SFTP server properties:

Property	Description
Enable SFTP Server	Whether to enable the SFTP server. When not enabled, the SFTP server cannot receive or send files. Default is disabled.
Port	Port number for the SFTP server. Default is 15002.
Local Address	Local IP address for the SFTP sever.
Enable SCP	Whether to use session control protocol (SCP) to create the connection. Default is disabled.
Idle Timeout	Number of seconds that the connection is idle before is closes. Default is 300.
Maximum Logins	Maximum number of users that can be logged in to the server concurrently. Default is 500.
Login Failure Delay	Delay between failed login attempts, in seconds. Default is 0.

Property	Description
Maximum Login Failures	Number of allowed login failures for a user. Default is 5.
Welcome Message	Message to show when the connection to the server is established.

Algorithms properties

Enable the following algorithm types in the **Algorithms** section of the **SFTP Server** tab:

- Cipher algorithms
- Message Authentication Code (MAC) algorithms
- Compression algorithms
- Key exchange algorithms

When you configure the use of algorithms for SFTP file exchange, consider the following rules and guidelines:

- You can move algorithms between the **Available** and **Selected** lists. The File Integration Service applies the algorithms that are listed in the **Selected** list.
- If no algorithms are selected for an algorithm type, the File Integration Service applies all the algorithms that are listed in the **Available** list.
- The File Integration Service applies the algorithms in the order in which they are listed, from the top of the list to the bottom of the list. You can use the up and down arrows to change the order of the algorithms in list.

Host keys properties

The following table describes host keys properties:

Property	Description
RSA Key File Location	Location of the RSA host key file.
RSA Key Passphrase	Passphrase for the RSA key.
DSA Key File Location	Location of the DSA host key file.
DSA Key Passphrase	Passphrase for the DSA key.

Upload restrictions properties

You can specify the types of files to allow or deny in an SFTP file exchange. The following table describes the properties that control upload restrictions:

Property	Description
File Extension Filter Type	Whether to accept or deny the extensions in the File Extensions list. Use one of the following values: - Do Not Filter. Accept all file types. - Accept. Accept files with the extensions listed in the File Extensions property. - Deny. Do not allow files with the extensions listed in the File Extensions property. Default is **Do Not Filter**.
File Extensions	List of file extensions. Add the file extensions that correspond to the File Extension Filter Type. For example, to accept .csv and .txt files, in the File Extension Filter Type property, select **Accept**, and then add csv and txt to the list of file extensions. To add an extension to the list, type the extension in the text box and click **Add**. To remove an extension from the list, highlight the extension and click Delete.
File Extension Case Sensitive	Whether to factor case when you filter using the file extensions list. When enabled, files with extensions that do not match the case used in the File Extensions list cannot be uploaded. For example, if the file extension list includes csv but not CSV, files with the extension of csv can be uploaded but files with the extension of CSV cannot be uploaded. Default is disabled.
Allow Files with No Extension	Whether to allow files that do not include the extension in the file name. Default is disabled.
Allow Files with Extension	Whether to enable the file extension filter. When enabled, the file extension properties configured on this page determine which file types can be uploaded. Default is enabled.

Proxy server configuration properties

For each runtime environment that uses the File Integration Service, you can configure one or more proxy servers.

You configure proxy server properties on the **Proxy Server** tab of the **File Server for agent** page.

To add a proxy server, click **Add Proxy Configuration**, configure server settings, and click **Save**.

Note: You must also install the proxy server in the DMZ. For more information, see "Installing a file integration proxy server" on page 176.

Configure the following types of properties:

- General properties
- Service mappings properties, which associate internal file servers with the proxy server

General properties

The following table describes general proxy server properties:

Property	Description
Enabled	Whether or not the proxy server is enabled. Default is Yes.
Controller Address	External IP address of the server in the DMZ on which the proxy server listens for control connections from the organization file servers.
Controller Port	Port number for the server in the DMZ on which the proxy server listens for control connections from the organization file servers. Default is 9100.
Minimum Number of Threads	Minimum number of threads that are reserved for connections to the location where the proxy server is installed. Default is 10.
Maximum Number of Threads	Maximum number of simultaneous requests that the proxy server can handle. Default is 2000.
Thread Keep Alive Time	The number of seconds idle threads wait before terminating. Default is 60.

Service mappings properties

To configure service mappings for the proxy server, in the **Proxy Server Configuration** page, click **Add** next to **Service Mappings**, configure the mapping parameters, and click **OK**. You can add as many service mappings as required to associate internal file servers with the proxy server.

The following table describes the service mappings properties:

Property	Description
Label	Label of the mapping.
From Address	IP address of the proxy server.
From Port	Port number of the proxy server.
To Address	IP address of the internal file server.
To Port	Port number of the internal file server.
Load Balancer Rule	Name of the load balancing rule to use with the mapping. The name of the rule must be identical to the name that appears in the `proxy.xml` file, which is part of the proxy server installation. For more information, see "Installing a file integration proxy server" on page 176.

Installing a file integration proxy server

Install a file integration proxy server in the DMZ and configure server parameters. You can install the server on Windows and Linux operating systems.

Note: You must also enable the proxy server and configure server properties in Informatica Intelligent Cloud Services, in Administrator. For more information, see "Proxy server configuration properties" on page 174.

1. Copy the `fis-proxy-server.zip` file to the server in the DMZ.
2. Download Java 1.8 (OpenJDK or Oracle) and install it on the server in the DMZ.
3. From the `fis-proxy-server/bin` folder, edit one of the following files:
   - On a Windows operating system, edit `setenv.bat`.
   - On a Linux operating system, edit `setenv.sh`.

   a. Set JAVA_HOME to the JDK Home or the JRE home of Java 1.8.

   b. Set the folder path of fis-proxy-server to FIS_PROXY_HOME.
4. From the `fis-proxy-server/config` folder, edit the `proxy.xml` file and set values for the following variables:

Variable	Description
controllerAddress	External IP address of the server in the DMZ on which the proxy server listens for control connections from the organization file servers.
dataAddress	Internal IP address of the server in the DMZ on which the proxy server listens for data connections from the organization file servers.
proxyAddress	IP address of the server in the DMZ on which the proxy server listens for incoming connections.
forwardProxyLocalAddress	IP address of the server in the DMZ on which the proxy server establishes outbound connections to remote servers as a forward proxy.

   If required, change the port numbers.
5. To start the proxy server, run one of the following commands:
   - On a Windows operating system, run `fis-proxy.bat start`.
   - On a Linux operating system, run `fis-proxy.sh start`.
6. To stop the proxy server, run one of the following commands:
   - On a Windows operating system, run `fis-proxy.bat stop`.
   - On a Linux operating system, run `fis-proxy.sh stop`.

   The proxy server saves logs in the `fis-proxy-server/logs` folder.

Stopping and starting a file server

You can stop or start a File Integration Service file server on the **File Servers** page. Stop and start a file server after you make configuration changes.

Stopping and starting HTTPS, AS2, and SFTP servers

To stop or start an HTTPS, AS2, or an SFTP server, perform the following actions:

1. In Administrator, select **File Servers**.
2. On the **File Servers** tab, click the arrow next to the name of the Secure Agent that runs the server.
3. From the Actions menu, select one of the following options:
   - Start AS2 Server
   - Stop AS2 Server
   - Start HTTPS Server
   - Stop HTTPS Server
   - Start SFTP Server
   - Stop SFTP Server

   Informatica Intelligent Cloud Services adds an entry in the audit log to indicate the action.

Stopping and starting a proxy server

To stop or start a proxy server, perform the following actions:

1. In Administrator, select **File Servers**.
2. On the **File Servers** tab, select the Secure Agent that runs the File Integration Service on which to stop or start the proxy server.
3. On the **File Server for agent** page, select the **Proxy Server** tab.
4. From the Actions menu of the server to stop or start, select **Stop** or **Start**.

   Informatica Intelligent Cloud Services adds an entry in the audit log to indicate the action.

File server users

Create a user account for each remote partner that exchanges files with your organization. The user account enables the partner to exchange files with your server.

For each remote partner, configure the following types of properties:

- General properties such as user name, email address, and password.
- Server-specific properties for HTTPS, AS2, and SFTP servers.
- Folder permissions.

Note: File server user accounts are different from Informatica Intelligent Cloud Services user accounts. File server user accounts enable remote partner users to exchange files with your organization's file servers. Informatica Intelligent Cloud Services user accounts enable your users to access your Informatica Intelligent Cloud Services organization.

Configuring a file server user

Configure partner users so that partners can exchange files with your organization.

When you create a file server user, the user receives an email from Informatica Intelligent Cloud Services. If you choose to include a system-generated password when you configure the user, a generated password is included in the email.

1. In Administrator, click **File Servers** > **File Server Users**.
2. Click **Add User**.
3. Perform the following actions, and then click **Save**:
   - Enter general information about the user.
   - To enable the user to send files to AS2 servers in the organization, enable the AS2 protocol and configure AS2 settings.
   - To enable the user to exchange files with SFTP servers in the organization, enable the SFTP protocol and configure SFTP settings.
   - To enable the user to exchange files with HTTPS servers in the organization, enable the HTTPS protocol and configure HTTPS settings.
   - Add folder and file permissions for the user. By default, the user has all permissions on the default home directory.

File server user properties

Configure properties for file server users.

General properties

The following table describes general properties for the user:

Property	Description
Username	User name for the file server user.
Description	Description for the user.
Company name	Name of the company.
Email	The user's email address.
Password Generation	Whether to create a password for the user or enable the system to create a system-generated password. Passwords must include the following characteristics: - Must be at least eight characters long. - Must have at least one upper case letter. - Must have at least one digit. - Must have at least one of the following special characters: `@ $ ! & * ~ - _`

AS2 server properties

The following table describes AS2 server properties for the file server user:

Property	Description
Enable AS2 Protocol	Whether or not the AS2 protocol is enabled. Disable when you do not want the AS2 server to receive files. Default is enabled.
Authentication Type	Whether to require a **Password**, **Certificate**, **Both**, or **Either**. If **Password** is used for authentication, the password generation that you defined in the **General** tab is used. If **Certificate** is used for authentication, the Client Authentication setting for the AS2 server must be set to Optional or Required.
SHA1 Fingerprint	If **Certificate**, **Both**, or **Either** is used for authentication, enter the SHA1 fingerprint of the partner's certificate. You can copy the SHA1 fingerprint of the certificate from the trust store.
AS2 ID	The AS2 ID of the partner user.
Signature Certificate Alias	Private key alias to use to sign the message. The private key is located in the default private key store.
Default Upload Folder	The location where AS2 files are saved when received. The default location is the default home directory for the user. If blank, the files are saved to the home directory. For more information, see "Global settings" on page 182.

SFTP server properties

The following table describes SFTP server properties for the file server user:

Property	Description
Enable SFTP Protocol	Whether or not the SFTP protocol is enabled. Disable when you do not want the SFTP server to send and receive files. Default is enabled.
Authentication Type	Whether to require a **Password**, **Public Key**, **Both** or **Either**. If a password is used for authentication, the password generation that you defined in the **General** tab is used. If a **Public Key** is used for authentication, you must place the key on the Secure Agent.
Public Key Location	Absolute path to the location of the public key on the Secure Agent. Applies when the **Public Key**, **Both**, or **Either** is used for authentication.

HTTPS server properties

The following table describes HTTPS server properties for the file server user:

Property	Description
Enable HTTPS Protocol	Whether or not the HTTPS protocol is enabled. Disable when you do not want the HTTPS server to receive files. Default is enabled.
Authentication Type	Whether to require a **Password**, **Certificate**, or **Either**. If a password is used for authentication, the password generation that you defined in the **General** tab is used. If **Certificate** is used for authentication, the Client Authentication setting for the HTTPS server must be set to Optional or Required.
SHA1 Fingerprint	If **Certificate** or **Either** is used for authentication, enter the SHA1 fingerprint of the partner's certificate. You can copy the SHA1 fingerprint of the certificate from the trust store.

Folder permissions properties

By default, a home directory and a user name are created for the user under the default home directory that is defined in the file servers **Global Settings** tab, and the user has all permissions on their home directory. You can edit the user's home directory to be in a different location.

To add permissions to other folders and files, click **Add** and define the permissions.

The following table describes folder permissions properties for the user:

Property	Description
Alias	An alias for the folder or file to which you grant permissions. The alias appears under the **Name** column in the **File Server User** page.
Path	Path to the folder or file to which you grant the user permissions.
Type	Determines whether the permissions are on a folder or on a file.
Folder Permissions	The user's permissions on the folder.
File Permissions	The user's permissions on the file.
Disk Space Restriction	Whether to restrict the disk space that the user can use on the folder, and, if yes, the allowed space on the disk. Applies to folder permissions.

Deleting a file server user

You might want to delete a file server user if the user no longer works with your organization..

1. In Administrator, select **File Servers** > **File Server Users**.
2. In the row that contains the user name, click Actions and select **Delete User**.

File transfer tasks

File transfer tasks are associated with the inbound and outbound process of the partner.

The **File Servers** page lists all the runtime environments and Secure Agents in your organization that can use the File Integration Service.

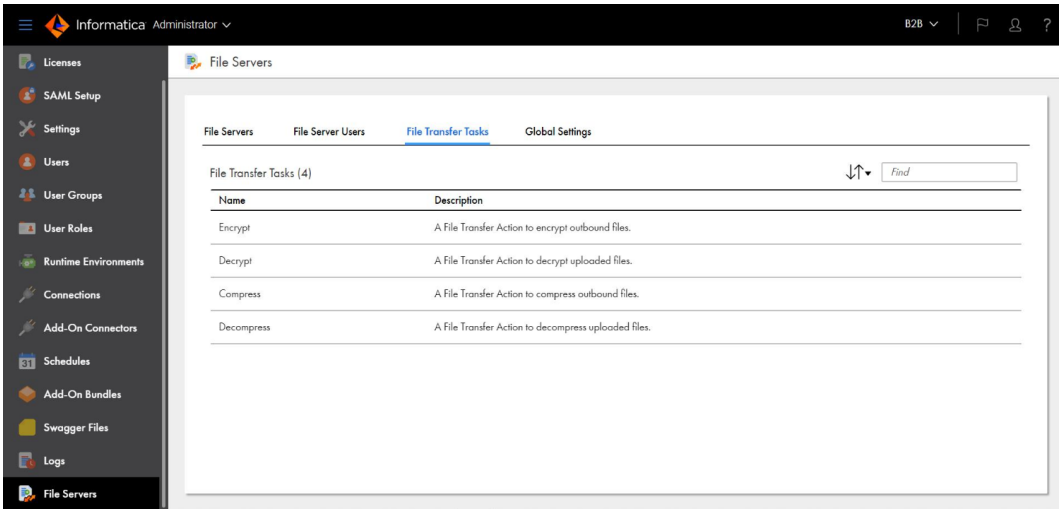

The **File Transfer Tasks** tab lists predefined file transfer tasks that can be used to run actions when files are received or sent to the file servers. The tab lists the projects in a read-only mode.

The **File Transfer Tasks** tab includes the following predefined file transfer tasks:

Name	Description
Encrypt	A file transfer task that uses PGP to encrypt outbound files when transferring them from the source location to the home directory of the file server user.
Decrypt	A file transfer task that uses PGP to decrypt uploaded files when transferring them from the home directory of file server user to the target location.
Compress	A file transfer task that compresses outbound files when transferring them from the source location to the home directory of file server user. You can choose one of the following compression methods: Zip, Tar, or Gzip.
Decompress	A file transfer task that decompresses uploaded files when transferring them from the home directory of file server user to the target location. You can choose one of the following decompression methods: Unzip, Untar, or Gunzip.

You can use the REST APIs to run the pre-defined tasks. You can also run the tasks using B2B Gateway.

For more information, see *REST API Reference*.

Global settings

Configure properties that apply to all file servers configured for file transfer.

Folder settings

Configure the **Default Home Directory** property to specify the default directory where all files received from remote servers are stored. User-specific home directories are created under the global home directory.

Note: Whenever you change the value of the default home directory you must stop and start the file servers.

SMTP server settings

The following table lists the SMTP server settings that apply to all remote file servers:

Property	Description
Host	Host name for the SMTP configuration to use for an email type MDN.
Port	Port on which the SMTP is running.
Username	User name to connect to the SMTP server.
Password	Password for the SMTP user.
Connection Type	Type of SMTP connection. Select one of the following values: - normal - implicitSSL - explicitSSL Default is normal.
From Email	Email address from which the email MDN is sent.
From Name	Name shown in the email.

PGP settings

Configure the **Public Key Ring** and **Secret Key Ring** to specify the directory where the public and secret key are stored. If the path is not specified, the default PGP key ring path is used.

Note: The PGP Command Line Interface (PGP-CLI) helps to edit the configuration properties file when you have multiple secure agents. You must restart the FIS application for the PGP setting changes to reflect in the `pgp-configuration.properties` file. The configuration file is present in the `conf` folder of the PGP client that is bundled with the FIS package.

CHAPTER 19

Troubleshooting

Use the following sections to troubleshoot errors in Administrator.

For a list of common error messages and possible solutions, see the Informatica Cloud Community article, "Troubleshooting: Common Error Messages".

Troubleshooting a Secure Agent

I installed the Secure Agent, but I want to install another on a different machine. How do I do that?

On the new machine, use your login to connect to Data Integration. Then, download and install the Secure Agent.

Secure Agent errors

I started the Secure Agent, but the status is inactive.

The Secure Agent might take a few minutes to start. The status refreshes every 5 seconds. If the Secure Agent does not become active, complete the following tasks:

- If your organization uses a proxy server to access the internet, verify that the proxy settings are set correctly.
- View the details in infaagent.log in the directory where you installed the Secure Agent.

The Secure Agent did not install or start correctly.

If the Secure Agent does not install or start correctly, complete the following tasks:

1. View the installation details in infaagent.log in the directory where you installed the Secure Agent.
2. View the application logs in the Event Viewer for a Secure Agent that runs on Windows.

One of my services shows an error status after I restarted the service successfully.

If a service fails with an error status, the error status for the service might continue to display in the Agent Service Details after the service starts up successfully. The error stays on the page until an internal job that cleans up obsolete messages runs. You can ignore the error.

I am trying to uninstall the Secure Agent, but the Secure Agent status still shows "Up and Running."

When you uninstall the Secure Agent without first stopping the Secure Agent, the Agent Core and other services might continue to run for several minutes. To avoid this issue, stop the Secure Agent before you uninstall it.

Troubleshooting an elastic cluster on AWS

Why did the elastic cluster fail to start?

To find out why the elastic cluster failed to start, use the `ccs-operation.log` file in the following directory on the Secure Agent machine:

```
<Secure Agent installation directory>/apps/At_Scale_Server/<version>/ccs_home/
```

The following table lists some reasons why a cluster might fail to start:

Reason	Possible Cause
kops failed to update the cluster.	The VPC limit was reached on your AWS account.
The master node failed to start.	The master instance type isn't supported in the specified region or availability zone or in your AWS account.
All worker nodes failed to start.	The worker instance type isn't supported in the specified region or availability zone or in your AWS account.
The Kubernetes API server failed to start.	The user-defined master role encountered an error.

When a cluster fails to start due to at least one of these reasons, the `ccs-operation.log` file displays a BadClusterConfigException.

For example, you might see the following error:

```
2019-06-27 00:50:02.012 [T:000060] SEVERE : [CCS_10500] [Operation of <cluster instance
ID>: start_cluster-<cluster instance ID>]:
com.informatica.cloud.service.ccs.exception.BadClusterConfigException: [[CCS_10207] The
cluster configuration for cluster [<cluster instance ID>] is incorrect due to the
following error: [No [Master] node has been created on the cluster. Verify that the
instance type is supported.]. The Cluster Computing System will stop the cluster soon.]
```

If the cluster encounters a BadClusterConfigException, the agent immediately stops the cluster to avoid incurring additional resource costs and to avoid potential resource leaks. The agent does not attempt to recover the cluster until the configuration error is resolved.

I looked at the ccs-operation.log file to troubleshoot the elastic cluster, but there wasn't enough information. Where else can I look?

You can look at the `cluster-operation` logs that are dedicated to the instance of the elastic cluster. When an external command set begins running, the `ccs-operation` log displays the path to the `cluster-operation` logs.

For example:

```
2020-06-15 21:22:36.094 [reqid:] [T:000057] INFO   :
c.i.c.s.c.ClusterComputingService        [CCS_10400] Starting to run command set
[<command set>] which contains the following commands: [
    <commands> ;
]. The execution log can be found in the following location: [/data2/home/cldagnt/
SystemAgent/apps/At_Scale_Server/35.0.1.1/ccs_home/3xukm9iqp5zeahyrb7rqoz.k8s.local/infa/
cluster-operation.log].
```

The specified folder contains all `cluster-operation` logs that belong to the instance of the cluster. You can use the logs to view the full `stdout` and `stderr` output streams of the command set.

The number in the log name indicates the log's generation and each `cluster-operation` log is at most 10 MB. For example, if the cluster instance generated 38 MB of log messages while running external commands, the folder contains four `cluster-operation` logs. The latest log has 0 in the file name and the oldest log has 3 in the file name. You can view the messages in the `cluster-operation0.log` file to view the latest errors.

If you set the log level for the Elastic Server to DEBUG, the `ccs-operation` log shows the same level of detail as the `cluster-operation` logs.

I ran a job to start the elastic cluster, but the VPC limit was reached.

When you do not specify a VPC in the elastic configuration for a cluster, the Secure Agent creates a new VPC on your AWS account. Because the number of VPCs on your AWS account is limited for each region, you might reach the VPC limit.

If you reach the VPC limit, edit the elastic configuration and perform one of the following tasks:

- Provide a different region.
- Remove the availability zones. Then, provide an existing VPC and specific subnets within the VPC for the cluster to use.

Any cloud resources that were provisioned for the cluster will be reused when the cluster starts in the new region or the existing VPC. For example, the Secure Agent might have provisioned Amazon EBS volumes before it received an error for the VPC limit. The EBS volumes are not deleted, but they are reused during the next startup attempt.

I ran a job to start the elastic cluster, but the cluster failed to be created with the following error:

```
Failed to create cluster [<cluster instance ID>] due to the following error:
[[CCS_10302] Failed to invoke AWS SDK API due to the following error: [Access Denied
(Service: Amazon S3; Status Code: 403; Error Code: AccessDenied; Request ID: <request
ID>; S3 Extended Request ID: <S3 extended request ID>)].].]
```

The Secure Agent failed to create the elastic cluster because Amazon S3 rejected the agent's request.

Make sure that the S3 bucket policies do not require clients to send requests that contain an encryption header.

How do I troubleshoot a Kubernetes API Server that failed to start?

If the Kubernetes API Server fails to start, the elastic cluster fails to start. To troubleshoot the failure, use the Kubernetes API Server logs instead.

To find the Kubernetes API Server logs, complete the following tasks:

1. Connect to the master node from the Secure Agent machine.
2. On the master node, locate the Kubernetes API Server log files in the directory `/var/log/`.

I updated the staging location for the elastic cluster. Now elastic mappings fail with the following error:

```
Error while executing mapping. ExecutionId '<execution ID>'. Cause: [Failed to start
cluster for [01000D25000000000005]. Error reported while starting cluster [Cannot apply
cluster operation START because the cluster is in an error state.].].
```

Mappings fail with this error when you change the permissions to the staging location before you change the S3 staging location in the elastic configuration.

If you plan to update the staging location, you must first change the S3 staging location in the elastic configuration and then change the permissions to the staging location on AWS. If you used role-based security, you must also change the permissions to the staging location on the Secure Agent machine.

To fix the error, perform the following tasks:

1. Revert the changes to the permissions for the staging location.
2. Edit the elastic configuration to revert the S3 staging location.
3. Stop the cluster when you save the configuration.
4. Update the S3 staging location in the configuration, and then change the permissions to the staging location on AWS.

I updated the staging location for the elastic cluster. Now the following error message appears in the agent job log:

```
Could not find or load main class com.informatica.compiler.InfaSparkMain
```

The error message appears when cluster nodes fail to download Spark binaries from the staging location due to access permissions.

Verify access permissions for the staging location based on the type of connectors that the job uses:

Connectors with direct access to Amazon data sources

If you use credential-based security for elastic jobs, make sure that the credentials in the Amazon S3 V2 and Amazon Redshift V2 connections can be used to access the staging location.

If you use role-based security for elastic jobs, make sure that the elastic cluster and the staging location exist under the same AWS account.

Connectors without direct access to Amazon data sources

If you use a user-defined worker role, make sure that the worker role can access both the staging location and the data sources in the elastic job.

If you use the default worker role, make sure that the Secure Agent role can access both the staging location and the data sources in the elastic job.

What should I do if the status of the elastic cluster is Unknown?

When the cluster status is Unknown, first verify that the Secure Agent is running. If the agent is not running, enable the agent and check whether the cluster starts running.

If the cluster does not start running, an administrator can run the command to list clusters. If the command output returns the cluster state as partial or in-use, the administrator can run the command to delete the cluster.

For more information about the commands, see *Data Integration Elastic Administration* in the Administrator help.

I restarted the Secure Agent machine and now the status of the elastic cluster is Error.

Make sure that the Secure Agent machine and the Secure Agent are running. Then, stop the elastic cluster in Monitor. In an AWS environment, the cluster might take 3 to 4 minutes to stop. After the cluster stops, you can run an elastic job to start the cluster again.

How do I find the initialization script logs for the nodes where the init script failed?

To find the init script logs, complete the following tasks:

1. Locate the `ccs-operation.log` file in the following directory on the Secure Agent machine:

   ```
   <Secure Agent installation directory>/apps/At_Scale_Server/<version>/ccs_home/
   ```

2. In the `ccs-operation.log` file, find a message that is similar to the following message:

   ```
   Failed to run the init script for cluster [<cluster instance ID>] on the following
   nodes: [<cluster node IDs]. Review the log in the following S3 file path: [<cloud
   platform location>].
   ```

3. Navigate to the cloud platform location that is provided in the message.
4. Match the cluster node IDs to the init script log file names for the nodes where the init script failed.

How are the resource requirements calculated in the following error message for an elastic cluster?

```
2019-04-26T19:04:11.762+00:00 <Thread-16> SEVERE: java.lang.RuntimeException:
[java.lang.RuntimeException: The Cluster Computing System rejected the Spark task
[InfaSpark0] due to the following error: [[CCS_10252] Cluster
[6bjwune8v4bkt3vneokii9.k8s.local] doesn't have enough resources to run the application
[spark--infaspark0e6674748-b038-4e39-a2a9-3fd49e63f289infaspark0-driver] which requires
a minimum resource of [(KB memory, mCPU)]. The cluster must have enough nodes, and each
node must have at least [(KB memory, mCPU)] to run this job.].]
```

The first resource requirement is the total number of resources that are required by the Spark driver and the Spark executor.

The second resource requirement is calculated based on the minimum resource requirements on each worker node to run a minimum of one Spark process.

The resources are calculated using the following formulas:

```
Memory: MAX(driver_memory, executor_memory)
CPU: MAX(driver_CPU, executor_CPU)
```

The Spark process can be either a Spark driver process or a Spark executor process. The cluster must have two nodes where each node fulfills the minimum requirements to run either the driver or the executor, or the cluster must have one node with enough resources to run both the driver and the executor.

Note: The resource requirements for the driver and executor depend on how you configure the following advanced session properties in the mapping task:

spark.driver.memory

spark.executor.memory

spark.executor.cores

For more information about minimum resource requirements, see *Data Integration Elastic Administration* in the Administrator help.

Is there anything I should do before I use a custom AMI to create cluster nodes?

If you use a custom AMI (Amazon machine image) to create cluster nodes, make sure that the AMI contains an installation of the AWS CLI.

The Secure Agent uses the AWS CLI to propagate tags to Amazon resources and to aggregate logs. The cluster nodes also use the AWS CLI to run initialization scripts.

For information about how to use a custom AMI, contact Informatica Global Customer Support.

My VPC has requirements to restrict internet traffic. Can I configure an elastic cluster to comply with these requirements?

By default, an elastic cluster uses an internet-facing load balancer to route traffic over the internet.

To restrict internet traffic, you can configure an elastic cluster to use an internal load balancer.

To use an internal load balancer, perform the following tasks:

1. To enable the internal load balancer, contact Informatica Global Customer Support.
2. Specify a VPC and subnets in the elastic configuration.
3. Make sure that the subnets use a NAT gateway so that cluster dependencies can be downloaded from the internet.

For more information about internet-facing and internal load balancers, refer to the AWS documentation.

Troubleshooting an elastic cluster on Microsoft Azure

After I set up staging and log locations on Blob Storage, the elastic mapping fails with the following error message in the session log:

```
20-02-11T00:52:43.273+00:00 <WorkflowExecutorThread20> INFO: [LDTM_0075] Total time to
perform the LDTM operation: 84,962 ms
2020-02-11T00:52:43.305+00:00 <InfaDisnextHadoopMappingExecutor-3-64> SEVERE:
java.lang.RuntimeException: java.lang.RuntimeException: java.lang.RuntimeException:
Failed to upload the local file in the path [/mnt/resource/informatica/secureagent/
apps/At_Scale_Server/33.0.1.1/metadata/0100edc7-f043-43f7-a5e1-a39f0774c2c7InfaSpark0/
submit_InfaSpark0_staticCode.jar] to the following shared storage location: [<Blob
Storage location>] due to the following error: [java.lang.RuntimeException:
[org.apache.hadoop.fs.azure.AzureException:
com.microsoft.azure.storage.StorageException: The account being accessed does not
support http.]].
2020-02-11T00:52:43.306+00:00 <InfaDisnextHadoopMappingExecutor-3-64> INFO: Spark
Mapping Ended with state: Failed
```

The error appears because Blob Storage requires requests to be made over HTTPS. To resolve the error, use the Azure portal to disable the `Secure transfer required` option for the storage account that holds the staging and log locations.

What should I do if the status of the elastic cluster is Unknown?

When the cluster status is Unknown, first verify that the Secure Agent is running. If the agent is not running, enable the agent and check whether the cluster starts running.

If the cluster does not start running, an administrator can run the command to list clusters. If the command output returns the cluster state as partial or in-use, the administrator can run the command to delete the cluster.

For more information about the commands, see *Data Integration Elastic Administration* in the Administrator help.

I restarted the Secure Agent machine and now the status of the elastic cluster is Error.

Make sure that the Secure Agent machine and the Secure Agent are running. Then, stop the elastic cluster in Monitor. In an Azure environment, the cluster might take 10 minutes to stop. After the cluster stops, you can run an elastic job to start the cluster again.

How do I find the initialization script logs for the nodes where the init script failed?

To find the init script logs, complete the following tasks:

1. Locate the `ccs-operation.log` file in the following directory on the Secure Agent machine:

   ```
   <Secure Agent installation directory>/apps/At_Scale_Server/<version>/ccs_home/
   ```

2. In the `ccs-operation.log` file, find a message that is similar to the following message:

   ```
   Failed to run the init script for cluster [<cluster instance ID>] on the following
   nodes: [<cluster node IDs]. Review the log in the following S3 file path: [<cloud
   platform location>].
   ```

3. Navigate to the cloud platform location that is provided in the message.
4. Match the cluster node IDs to the init script log file names for the nodes where the init script failed.

The init script failed with the following standard error on some nodes in the elastic cluster:

```
Created symlink from /etc/systemd/system/apt-daily.service to /dev/null.
Created symlink from /etc/systemd/system/apt-daily-upgrade.service to /dev/null.
Removed symlink /etc/systemd/system/timers.target.wants/apt-daily.timer.
Removed symlink /etc/systemd/system/timers.target.wants/apt-daily-upgrade.timer.
E: Could not get lock /var/lib/dpkg/lock-frontend - open (11: Resource temporarily
unavailable)
E: Unable to acquire the dpkg frontend lock (/var/lib/dpkg/lock-frontend), is another
process using it?
```

The init script failed because the node was running an internal process at the same time as the init script. If you continue to see the error, wait for the internal process to complete by placing a sleep command for the required duration in your init script.

For example, you might use a sleep command as follows:

```
#!/bin/sh

while(sudo lsof /var/lib/dpkg/lock-frontend)
do
echo "Sleeping 10s"
sleep 10
done

sudo apt-get -y update
sudo apt-get install -y expect
```

I looked at the ccs-operation.log file to troubleshoot the elastic cluster, but there wasn't enough information. Where else can I look?

You can look at the `cluster-operation` logs that are dedicated to the instance of the elastic cluster. When an external command set begins running, the `ccs-operation` log displays the path to the `cluster-operation` logs.

For example:

```
2020-06-15 21:22:36.094 [reqid:] [T:000057] INFO   :
c.i.c.s.c.ClusterComputingService        [CCS_10400] Starting to run command set
[<command set>] which contains the following commands: [
    <commands> ;
]. The execution log can be found in the following location: [/data2/home/cldagnt/
```

```
SystemAgent/apps/At_Scale_Server/35.0.1.1/ccs_home/3xukm9iqp5zeahyrb7rqoz.k8s.local/infa/
cluster-operation.log].
```

The specified folder contains all `cluster-operation` logs that belong to the instance of the cluster. You can use the logs to view the full `stdout` and `stderr` output streams of the command set.

The number in the log name indicates the log's generation and each `cluster-operation` log is at most 10 MB. For example, if the cluster instance generated 38 MB of log messages while running external commands, the folder contains four `cluster-operation` logs. The latest log has 0 in the file name and the oldest log has 3 in the file name. You can view the messages in the `cluster-operation0.log` file to view the latest errors.

If you set the log level for the Elastic Server to DEBUG, the `ccs-operation` log shows the same level of detail as the `cluster-operation` logs.

How are the resource requirements calculated in the following error message for an elastic cluster?

```
2019-04-26T19:04:11.762+00:00 <Thread-16> SEVERE: java.lang.RuntimeException:
[java.lang.RuntimeException: The Cluster Computing System rejected the Spark task
[InfaSpark0] due to the following error: [[CCS_10252] Cluster
[6bjwune8v4bkt3vneokii9.k8s.local] doesn't have enough resources to run the application
[spark--infaspark0e6674748-b038-4e39-a2a9-3fd49e63f289infaspark0-driver] which requires
a minimum resource of [(KB memory, mCPU)]. The cluster must have enough nodes, and each
node must have at least [(KB memory, mCPU)] to run this job.].]
```

The first resource requirement is the total number of resources that are required by the Spark driver and the Spark executor.

The second resource requirement is calculated based on the minimum resource requirements on each worker node to run a minimum of one Spark process.

The resources are calculated using the following formulas:

```
Memory: MAX(driver_memory, executor_memory)
CPU: MAX(driver_CPU, executor_CPU)
```

The Spark process can be either a Spark driver process or a Spark executor process. The cluster must have two nodes where each node fulfills the minimum requirements to run either the driver or the executor, or the cluster must have one node with enough resources to run both the driver and the executor.

Note: The resource requirements for the driver and executor depend on how you configure the following advanced session properties in the mapping task:

spark.driver.memory

spark.executor.memory

spark.executor.cores

For more information about minimum resource requirements, see *Data Integration Elastic Administration* in the Administrator help.

Troubleshooting scheduled tasks

The task does not run at the scheduled time.

A task does not run at the scheduled time if another instance of it is already running when the schedule tries to start the task. For example, you schedule a task to run every 5 minutes. The first task starts at 12 p.m., but

does not complete until 12:06 p.m. The second instance of the task does not run at 12:05 p.m. because the first instance has not completed. Data Integration starts the next task at 12:10 p.m.

To resolve this issue, change the schedule to allow the task to complete before starting the next task run.

Troubleshooting security

I received the following security violation error:

```
There may have been a security violation while accessing the site. Verify that there are
no malicious scripts running in your browser. This error also appears when you submit
the form multiple times through a browser reload.
```

This error appears when you click an option on a page while the page is still loading from a previous click. Click the Here link to return to Data Integration.

When I try to view the details about an object, such as a connection or replication task, the Object Not Found page displays.

The object was recently deleted. The Object Not Found page appears when an object no longer exists. Refresh the page to display current objects.

When I try to perform a task, the Access Denied page displays.

The Access Denied page displays when you try to perform a task that is not allowed for your user account. You might not have the appropriate role or asset permissions to perform the task. If you need to perform the task, ask your organization administrator to review your user account.

INDEX

R

S

T

U

W