



Informatica® Operational Insights
August 2024

Operational Insights

Informatica Operational Insights Operational Insights
August 2024

© Copyright Informatica LLC 2017, 2024

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

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

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

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

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

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

Publication Date: 2024-08-07

Table of Contents

| | |
|--|-----------|
| Preface | 7 |
| Part I: Introducing Operational Insights | 8 |
| Chapter 1: Operational Insights Overview | 9 |
| Configure alerts | 10 |
| Understanding collectors | 11 |
| Monitor services and applications with Operational Insights | 13 |
| Informatica Resources | 13 |
| Informatica Documentation | 13 |
| Informatica Intelligent Cloud Services web site | 13 |
| Informatica Intelligent Cloud Services Communities | 13 |
| Informatica Intelligent Cloud Services Marketplace | 14 |
| Data Integration connector documentation | 14 |
| Informatica Knowledge Base | 14 |
| Informatica Intelligent Cloud Services Trust Center | 14 |
| Informatica Global Customer Support | 14 |
| Part II: Monitoring Informatica Intelligent Cloud Services | 15 |
| Chapter 2: Install and configure a Secure Agent | 16 |
| Secure Agent prerequisites | 16 |
| Installing in Windows | 17 |
| Secure Agent requirements on Windows | 17 |
| Downloading and installing the Secure Agent on Windows | 18 |
| Configuring the firewall | 19 |
| Configuring the proxy settings in Windows | 20 |
| Configuring a login for a Windows Secure Agent service | 21 |
| Uninstalling the Secure Agent in Windows | 21 |
| Installing in Linux | 22 |
| Secure Agent requirements on Linux | 22 |
| Downloading and installing the Secure Agent on Linux | 23 |
| Configuring the firewall | 24 |
| Configuring the proxy settings in Linux | 25 |
| Uninstalling the Secure Agent in Linux | 26 |
| Chapter 3: Monitor Informatica Intelligent Cloud Services infrastructure | 27 |
| Adding a Secure Agent or domain to a map location | 28 |
| Monitoring infrastructure health | 28 |
| Monitoring Secure Agents | 29 |

| | |
|---|-----------|
| Monitoring the OI Data Collector service. | 30 |
| Zooming in on graph details. | 30 |
| Infrastructure alerts. | 31 |
| Configuring infrastructure alerts. | 32 |
| Use alert scripts. | 33 |
| Configuring alerts for Secure Agent services. | 35 |
| Chapter 4: Monitor Informatica Intelligent Cloud Services Data Integration. | 36 |
| View Data Integration analytics. | 37 |
| View Data Integration jobs. | 39 |
| Configure Jobs page settings. | 41 |
| Exporting jobs data. | 42 |
| View details for a specific job. | 43 |
| View job history for a specific asset. | 44 |
| View Data Integration connections. | 46 |
| View Data Integration connection events. | 47 |
| View scheduled jobs. | 48 |
| Data Integration alerts. | 48 |
| Configuring alerts for Data Integration jobs. | 49 |
| Chapter 5: Monitor Informatica Intelligent Cloud Services Application Integration. | 50 |
| View overall usage and health of Application Integration assets. | 50 |
| View Application Integration process reports. | 57 |
| Monitor usage of Application Integration assets. | 58 |
| View incoming API calls. | 58 |
| View Application Integration process runs. | 60 |
| View Application Integration connection calls. | 63 |
| View API transactions against licensing limit. | 66 |
| Chapter 6: Monitor Informatica Intelligent Cloud Services Data Profiling. | 68 |
| View Data Profiling service jobs. | 69 |
| View details for a specific job in data profiling. | 70 |
| Chapter 7: Monitor Informatica Intelligent Cloud Services Data Ingestion and Replication | 72 |
| Monitoring your ingestion and replication jobs. | 73 |
| Monitoring all ingestion and replication jobs. | 74 |
| Job properties. | 77 |
| Viewing details for an ingestion and replication job. | 79 |
| Application ingestion and replication job details. | 79 |
| Database ingestion and replication job details. | 85 |
| File ingestion and replication job details. | 92 |

| | |
|--|------------|
| Streaming ingestion and replication job details. | 95 |
| Data Ingestion and Replication alerts. | 97 |
| Configuring alerts for Data Ingestion and Replication jobs. | 98 |
| Chapter 8: Monitor MDM SaaS. | 99 |
| Prerequisites to monitor MDM SaaS. | 99 |
| Monitor usage statistics for MDM SaaS. | 99 |
| Monitor MDM SaaS jobs. | 100 |
| View job summaries. | 100 |
| View job instances and schedules. | 101 |
| View key job metrics. | 102 |
| Ingress job metrics. | 103 |
| Egress job metrics. | 103 |
| Match job metrics. | 104 |
| Merge job metrics. | 104 |
| View job performance. | 104 |
| Export jobs. | 105 |
| Compare jobs. | 105 |
| Part III: Monitoring on-premises applications. | 107 |
| Chapter 9: Register and manage domains. | 108 |
| Enable the monitoring Model Repository Service. | 108 |
| Configuring the domain connection. | 109 |
| Entering the domain details. | 110 |
| Configure the Domain Configuration Collector. | 111 |
| Configure the Collector Schedule. | 111 |
| Configure the Domain Health Statistics Collector. | 111 |
| Configure the Collector Schedule. | 112 |
| Configure the Domain Resource Usage Statistics Collector. | 112 |
| Collecting historical data. | 112 |
| Connecting to the Monitoring Statistics Model repository. | 113 |
| Configure the Collector Schedule. | 114 |
| Configure the PowerCenter Repository Collector. | 114 |
| Collecting historical data. | 114 |
| Adding a PowerCenter repository. | 114 |
| Configure the Collector Schedule. | 116 |
| Configure the Data Engineering Integration collector. | 116 |
| Collecting historical data. | 116 |
| Selecting the cluster configuration. | 117 |
| Configure the Collector Schedule. | 117 |
| Connecting to a cluster secured using Kerberos authentication. | 117 |
| Configure the Data Quality Collector. | 118 |

| | |
|---|------------|
| Configure the Collector Schedule. | 118 |
| Finalize the on-boarding configuration. | 119 |
| Search for domains. | 119 |
| Editing or unregistering a domain. | 119 |
| Set the time zone. | 120 |
| | |
| Chapter 10: Monitor Data Engineering Integration domains | 121 |
| Viewing Data Engineering Integration job analytics. | 122 |
| Viewing job execution data. | 122 |
| Creating a Data Engineering Integration project. | 123 |
| | |
| Chapter 11: Monitor Data Quality domains | 124 |
| Viewing Data Quality job analytics. | 125 |
| Viewing job execution summary data. | 125 |
| Creating a Data Quality project. | 126 |
| | |
| Chapter 12: Monitor PowerCenter domains. | 127 |
| Viewing PowerCenter workflow analytics. | 128 |
| Viewing PowerCenter workflow execution data. | 128 |
| Viewing anomalous workflow run behavior. | 129 |
| Viewing recommendations. | 130 |
| Viewing the resource utilization heat map. | 131 |
| Viewing domain resource utilization. | 132 |
| Creating a PowerCenter project. | 133 |
| Using PowerCenter repository filters. | 134 |
| PowerCenter alerts. | 134 |
| | |
| Index. | 136 |

Preface

Use Informatica Intelligent Cloud ServicesSM Operational Insights to learn how to gain visibility into the performance and operational efficiency of your Informatica infrastructure, view analytics, and monitor jobs, assets, and connections. Learn how to use workflow and job execution metrics to troubleshoot failures, and how to review resource consumption analytics to predict when you might need to increase capacity or reallocate resources. Understand how to configure email notifications to alert you about problems within a domain or with a Secure Agent, and how to use insights in the form of recommendations and anomalous workflow detection to address potential issues.

Part I: Introducing Operational Insights

This part contains the following chapter:

- [Operational Insights Overview, 9](#)

CHAPTER 1

Operational Insights Overview

Operational Insights is a service in Informatica Intelligent Cloud Services that gives you visibility into the performance and operational efficiency of your Informatica infrastructure. Use Operational Insights to monitor your Informatica cloud services and Informatica on-premises products.

Operational Insights is aware of the runtime environments, Informatica Cloud Secure Agents, Secure Agent services, Informatica Intelligent Cloud Services, and domains in your infrastructure. Run-time statistics and asset configuration metadata is uploaded from assets to the service on configurable schedules, providing you with an accurate, up-to-date overview of your Informatica deployments.


The analytics-driven features in Operational Insights provide you with the following capabilities:

- Comprehensive monitoring statistics allow you to quickly assess the health of your Informatica assets.
- Data processing analytics based on job runs and data processing statistics enable you to assess usage of your investment in Informatica.
- Run-time job execution, workflow, and task metrics allow you to troubleshoot performance degradation and execution failures.
- Resource consumption analytics help you predict when you need to increase capacity or reallocate resources.
- Email notifications alert you about problems within a Data Engineering Integration, Data Quality, or PowerCenter domain or with a Secure Agent.
- Insights in the form of recommendations and anomalous workflow detection help you identify errors and abnormal behavior within a domain.

Configure alerts

You can configure Operational Insights to send alerts when an issue occurs in a service within the organization. Use the **Alerts** page to configure and manage alerts for Informatica assets.

The following image shows the Alerts page:



The **Alerts** page displays the following tabs:

Alerts Received

View information about all alert notifications received by a user for a service in the organization. You can choose to view alerts from the last 24 hours, the past week, or a custom date range. Default is 24 hours. You can also filter notifications based on the service. You can drill down on an alert name to view the alert details.

Infrastructure Alerts

Configure alerts for domains, Secure Agents, and Secure Agent services. You can configure alerts for the parent organization and sub-organizations. You can set the Secure Agent to send alert notification emails to recipients within the organization and outside the organization.

Note: During the first 30 days after migrating the organization to a new POD, you might receive email notifications informing you that the Secure Agent is down. The notifications refer to the Secure Agent that the organization used prior to the migration and can be ignored.

Data Integration Alerts

Configure and manage alerts for Data Integration jobs. You can configure alerts for jobs that are in a specific status or reach a threshold limit. You can also configure Operational Insights to restart or stop jobs that have reached a specific threshold limit.

Data Ingestion and Replication Alerts

Configure and manage alerts for application ingestion and replication and database ingestion and replication jobs. You can configure alerts for jobs that acquire a specific job state and apply the alerts across the entire selected organization or to a specific task asset.

PowerCenter Alerts

Configure and manage PowerCenter workflow alerts and CLAIRE alerts for anomalous or abnormal behavior in PowerCenter workflow instances. You can set alerts for issues that occur in workflows in the PowerCenter domain or in a project created in Operational Insights.

Understanding collectors

Collectors are components that run within the Secure Agent that communicates with an Informatica domain. You do not need to configure a collector to collect data from a specific asset. The Secure Agent is aware of all assets within the domain, and each enabled collector collects operational data and metadata from the relevant assets.

The collectors can collect data from all Informatica release 10.x domains. If needed, you can disable any collector except for the Domain Configuration Collector. You can also change the default collection schedule for each collector.

The collectors begin collecting and uploading data to Informatica Intelligent Cloud Services after you click **Save** as the final step in the domain registration process. Data is uploaded to Operational Insights at the time it is collected.

The following collectors are deployed with the Secure Agent:

| Collector | Description | Default Collection Frequency | Metadata Collected |
|--|--|------------------------------|--|
| Domain Configuration Collector | Collects and uploads configuration metadata for the domain and all domain assets, including nodes and services.* | Every 24 hours | <ul style="list-style-type: none"> - Domain details: Domain name, list of nodes in the domain, details for grids within the domain. - Node details: Name, HTTP port, logs directory, maximum processes, list of services running. - System configuration for each node: Operating system details, number of CPU cores and CPU speed, physical memory details. |
| Domain Health Statistics Collector | Collects and uploads availability statistics for domain assets, including nodes and application services.* | Every 5 minutes | <ul style="list-style-type: none"> - Availability statistics: Domains, nodes, services, and grids. |
| Domain Resource Usage Statistics Collector | <p>Collects and uploads CPU and memory consumption statistics for all nodes within the domain.</p> <p>The collector collects the statistics from the Model repository managed by the monitoring Model Repository Service specified in the Monitoring Configuration for the domain.</p> | Every 1 hour | <ul style="list-style-type: none"> - CPU utilization: Informatica processes and all processes running on each node. - Memory utilization: Informatica processes and all processes running on each node. |

| Collector | Description | Default Collection Frequency | Metadata Collected |
|--|---|------------------------------|---|
| Data Engineering Integration Collector | Collects and uploads statistics on Data Engineering Integration jobs run on Hadoop. You must specify the cluster configuration that contains the configuration information for each Hadoop cluster you want to collect statistics from. | Every 1 hour | <ul style="list-style-type: none"> - Hadoop cluster configuration: Number of clusters and nodes by Hadoop distribution. - Hadoop cluster resource usage: CPU and memory utilization for Informatica processes and all processes running on each cluster node. - Job execution statistics: Type of job, how many rows inserted/rejected, how many mappings failed/succeeded, volume of data processed, start and end time, Data Integration Service that submitted the job. - Runtime metrics: Number and type of unique mappings and workflows that ran on each cluster or on all clusters. |
| Data Quality Collector | Collects and uploads statistics on Data Quality jobs run on Hadoop. | Every 1 hour | <ul style="list-style-type: none"> - Job execution statistics: Type of job, how many rows inserted/rejected, how many mappings failed/succeeded, volume of data processed, start and end time, Data Integration Service that submitted the job. - Runtime metrics: Number and type of unique mappings and workflows that ran on each cluster or on all clusters. |
| PowerCenter Repository Collector | Collects and uploads runtime workflow and session metrics from PowerCenter repositories within the domain. You must provide the JDBC connection details for the repository database for each PowerCenter repository you want to collect statistics from. | Every 1 hour | <ul style="list-style-type: none"> - Workflow details: Workflow name, start and end time, status (succeeded, failed, etc.), PowerCenter Integration Service that submitted the workflow. - Session tasks: Task ID and type, start and end time, rows read, rows written, node the task ran on. - Repository folder details: Folder ID and name. |

**If the Informatica domain that you registered to monitor using Operational Insights has multiple gateway nodes and the domain fails over from the master gateway node to another, the Domain Configuration Collector and the Domain Health Statistics Collector internally switch the collection to other gateway nodes in the domain and continue to work seamlessly.*

Monitor services and applications with Operational Insights

Use Informatica Intelligent Cloud Services Operational Insights to monitor the following Informatica cloud services:

- Informatica environment
- Informatica Intelligent Cloud Services Application Integration
- Informatica Intelligent Cloud Services Data Integration
- Informatica Intelligent Cloud Services Data Quality
- Informatica Intelligent Cloud Services Data Profiling
- Informatica Intelligent Cloud Services Data Ingestion and Replication

Use Informatica Intelligent Cloud Services Operational Insights to monitor the following Informatica on-premises applications:

- Data Engineering Integration domains
- Data Quality domains
- PowerCenter domains

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, on the [Informatica Intelligent Cloud Services Status](#) page, click **SUBSCRIBE TO UPDATES**. You can 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 Global Support Center through the Informatica Network or by telephone.

To find online support resources on the Informatica Network, click **Contact Support** in the Informatica Intelligent Cloud Services Help menu to go to the **Cloud Support** page. The **Cloud Support** page includes system status information and community discussions. Log in to Informatica Network and click **Need Help** to find additional resources and to contact Informatica Global Customer Support through email.

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>.

Part II: Monitoring Informatica Intelligent Cloud Services

This part contains the following chapters:

- [Install and configure a Secure Agent, 16](#)
- [Monitor Informatica Intelligent Cloud Services infrastructure , 27](#)
- [Monitor Informatica Intelligent Cloud Services Data Integration, 36](#)
- [Monitor Informatica Intelligent Cloud Services Application Integration, 50](#)
- [Monitor Informatica Intelligent Cloud Services Data Profiling, 68](#)
- [Monitor Informatica Intelligent Cloud Services Data Ingestion and Replication , 72](#)
- [Monitor MDM SaaS, 99](#)

CHAPTER 2

Install and configure a Secure Agent

You must configure every domain that Operational Insights monitors to communicate with a Secure Agent. Secure Agent installation is required only for on-premises products. You do not install Secure Agents for Informatica Intelligent Cloud Services.

If the domain can access an existing Secure Agent, you can configure the domain to use the Secure Agent when you register the domain with Operational Insights. You can configure multiple domains within your organization to use the same Secure Agent.

If the domain doesn't have access to a Secure Agent, you must download and install a Secure Agent on a node within the domain. The Secure Agent must be installed on a machine that has access to the internet.

Note that if the Operational Insights service shuts down for more than 12 hours, the Secure Agent connection to the service times out. Manually restart the Secure Agent on the node to re-establish the connection.

Note: If you use Operational Insights to monitor other services in Informatica Intelligent Cloud Services, see the Administrator help for information about configuring a Secure Agent.

Secure Agent prerequisites

Perform the following tasks before you download and install the Secure Agent:

- On Windows, you must log in as a non-administrative user. On Linux, you must log in as a non-root user. If you install the Secure Agent as a user with administrative rights or with root access, the underlying Secure Agent database does not start.
- (Optional) If you want use JDBC and SAP connector third party libraries with the Process Server, copy the third party libraries to the following location:
`<Secure Agent Installation Directory>/apps/process-engine/ext`

Installing in 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. You can stop and restart the Secure Agent using the Secure Agent Manager or Windows Services. If you install the Secure Agent on a different volume than you use to run the installation program, you must start and stop the Secure Agent from Windows Services.

You can also use the Secure Agent Manager to check the Secure Agent status and configure proxy information. The Secure Agent works with BASIC, DIGEST, and NTLMv2 proxy authentication.

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 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:

- The Secure Agent machine 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 Knowledge Base.
- The Secure Agent machine has the Microsoft Visual C++ 2015 Redistributable.
- The Secure Agent machine has at least 4 CPU cores, 16 GB RAM, and at least 5 GB of free disk space.
- The Secure Agent machine is on a volume with at least 250GB disk space, with at least 5 GB free space or three times the size of the Secure Agent installation, whichever is greater.
- The account you use to install the Secure Agent has access to all remote directories that contain flat source or target files.
- No other Secure Agent is installed on the machine. If another Secure Agent is installed on the machine, uninstall it first.

For more information, see this KB article:

[Minimum requirements and best practices when installing Informatica Cloud Secure Agent.](#)

Setting 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 any other Secure Agent exists, you must uninstall it.

Tip: To verify the checksum of the Secure Agent installation program, use the agent REST API version 2 resource. For more information about the agent resource, see *REST API Reference*.

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.<agent core version>.exe`.
4. Run the installation program as an Administrator:
 - a. Specify the Secure Agent installation directory, and click **Next**.
 - b. Click **Install** to install the agent.

The **Cloud Secure Agent** dialog box 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.

Configuring the firewall

If your organization uses a protective firewall, include the Informatica Intelligent Cloud Services and Operational Insights domain name or IP address ranges for your region in the list of approved domain names

or IP addresses. You must also configure Secure Agents used by the domains Operational Insights monitors to use the approved IP address ranges.

You should also enable the port that the Secure Agent uses. The Secure Agent uses port 443 (HTTPS) to connect to the internet. Configure your firewall to allow traffic to pass over port 443.

You must add the domain name or IP address ranges required by both Informatica Intelligent Cloud Services and Operational Insights to your list of approved domain names or IP addresses.

The allowlists 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 domains and IP addresses to allowlist for Informatica Intelligent Cloud Services in the following KB article on Informatica Network:

<https://kb.informatica.com/faq/7/Pages/20/524982.aspx>

You can find the domains and IP addresses to allowlist for Operational Insights in the following KB article on Informatica Network:

<https://kb.informatica.com/faq/7/Pages/21/532624.aspx>

To configure a Secure Agent to use the approved IP address ranges, complete the steps below:

1. Add either the domain names or the IP address ranges for your region to your list of approved addresses.
2. Log in to Operational Insights.
3. Select a domain, and then click the **Details** tab.
4. Locate the name of the Secure Agent the domain uses in the Secure Agent Group property.
5. Click **Secure Agents** in the left hand navigation bar.
6. Select a Secure Agent, then click **Manage**.
The Details page for the Secure Agent opens in the Administrator application.
7. Click **Edit**.
8. Click the **+** symbol next to a property in the Custom Configuration section of the page to add a new custom property.
9. Select **Opsinsights Data Collector** from the Service menu, and then select **Opsinsights** from the Type menu.
10. Enter `useStaticIP` in the Name field, and then enter `true` in the Value field.
11. Click **Save**.

Configuring the proxy settings in 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.

Configuring 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 Secure Agent machine 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 in 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.

Note: Uninstalling the Secure Agent does not delete log files from the Secure Agent directory. If you want to reinstall a Secure Agent on the machine, you must delete all files and directories associated with the Secure Agent installation or reinstallation will fail. If you want to save the log files, copy them to a different directory, and then delete the Secure Agent installation directory.

Installing in 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 meets the minimum requirements.
2. Download the Secure Agent installer files.
3. Install and register the Secure Agent.

Consider the following guidelines:

- Create a specific user profile to install the Secure Agent with full access to all folders from the Secure Agent installation directory. Don't install the Secure Agent as the root user.
- You can't install more than one Secure Agent on the same machine under the same user account. Multiple agents may exist under different user accounts.
- Don't install the Secure Agent on any node within the Informatica domain.

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 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 Knowledge Base.

- Verify that the machine has at least 11 GB free disk space.

- Verify that the `libidn.x86_64` package is installed.

If the package isn't present, install it using the following command: `sudo yum install libidn.x86_64`

Note: The command to install the package might vary based on your Linux distribution.

- Verify that the `libidn.so.*` libraries are installed.

If the libraries aren't present, install them using the following commands:

- For 64-bit systems: `cd /usr/lib/x86_64-linux-gnu`

- For 32-bit systems: `cd /usr/lib/i386-linux-gnu`

After installing the libraries, create a symbolic link using the following command:

```
sudo ln -s libidn.so.12 libidn.so.11
```

- If you are installing the Secure Agent on RHEL 9, verify that the `libnsl` library is installed. If the library isn't present, install it using the following command: `sudo yum install libnsl`
Note: The command to install the package might vary based on your Linux distribution.
To verify whether `libnsl` is present, use one of the following commands: `ldconfig -p | grep libnsl` or `which libnsl`.
- 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, see this KB article:

[Minimum requirements and best practices when installing Informatica Cloud Secure Agent.](#)

Setting 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.

When you register the agent, it is added to its own Secure Agent group by default. You can add the agent to a different Secure Agent group.

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.

Tip: To verify the checksum of the Secure Agent installation program, use the agent REST API version 2 resource. For more information about the agent resource, see *REST API Reference*.

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.<agent core version>.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.

- 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
```

- When the installer completes, navigate to the following directory:

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

- 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.

- 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.
- To register the agent, in the `<Secure Agent installation directory>/apps/agentcore` directory, enter one of the following commands using your Informatica Intelligent Cloud Services user name and the token that you copied:

- To add the agent to its own Secure Agent group, use the following command:

```
./consoleAgentManager.sh configureToken <user name> <install token>
```

- To add the agent to an existing Secure Agent group, use the following command:

```
./consoleAgentManager.sh configureTokenWithRuntime <user name> <install token>  
<Secure Agent group name>
```

Note: If the command includes a Secure Agent group name that doesn't exist, the Secure Agent is not assigned to a group. Be sure to use a valid Secure Agent group name.

The following table lists the command options:

| Option | Description |
|-------------------------|--|
| User Name | Required. Informatica Intelligent Cloud Services user name of the user installing the Secure Agent. |
| Install Token | Required. The install token that you copied. |
| Secure Agent group name | Optional. Include when you want to add the agent to an existing Secure Agent group instead. If this option isn't included in the command, the agent will be in its own Secure Agent group. |

You can check the registration status of a Secure Agent using the following command:

```
./consoleAgentManager.sh isConfigured
```

Configuring the firewall

If your organization uses a protective firewall, include the Informatica Intelligent Cloud Services and Operational Insights domain name or IP address ranges for your region in the list of approved domain names or IP addresses. You must also configure Secure Agents used by the domains Operational Insights monitors to use the approved IP address ranges.

You should also enable the port that the Secure Agent uses. The Secure Agent uses port 443 (HTTPS) to connect to the internet. Configure your firewall to allow traffic to pass over port 443.

You must add the domain name or IP address ranges required by both Informatica Intelligent Cloud Services and Operational Insights to your list of approved domain names or IP addresses.

The allowlists 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 domains and IP addresses to allowlist for Informatica Intelligent Cloud Services in the following KB article on Informatica Network:

<https://kb.informatica.com/faq/7/Pages/20/524982.aspx>

You can find the domains and IP addresses to allowlist for Operational Insights in the following KB article on Informatica Network:

<https://kb.informatica.com/faq/7/Pages/21/532624.aspx>

To configure a Secure Agent to use the approved IP address ranges, complete the steps below:

1. Add either the domain names or the IP address ranges for your region to your list of approved addresses.
2. Log in to Operational Insights.
3. Select a domain, and then click the **Details** tab.
4. Locate the name of the Secure Agent the domain uses in the Secure Agent Group property.
5. Click **Secure Agents** in the left hand navigation bar.
6. Select a Secure Agent, then click **Manage**.
The Details page for the Secure Agent opens in the Administrator application.
7. Click **Edit**.
8. Click the **+** symbol next to a property in the Custom Configuration section of the page to add a new custom property.
9. Select **OpsInsights Data Collector** from the Service menu, and then select **OpsInsights** from the Type menu.
10. Enter `useStaticIP` in the Name field, and then enter `true` in the Value field.
11. Click **Save**.

Configuring the proxy settings in 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. The Secure Agent works with BASIC, DIGEST, and NTLMv2 proxy authentication.

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 in 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:

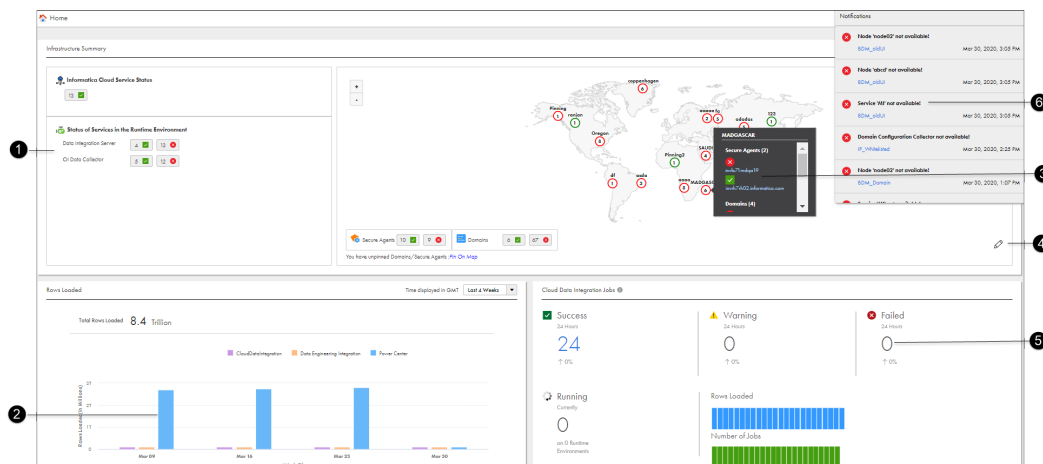
```
./infaagent shutdown
```

3. To uninstall the Secure Agent, run `rm -rf` on the directory where you installed the Secure Agent to remove Secure Agent files.

CHAPTER 3

Monitor Informatica Intelligent Cloud Services infrastructure

You can use Operational Insights to monitor your Informatica Intelligent Cloud Services infrastructure. Use the **Home** page to assess the overall usage and health of your Informatica assets. Parent organization users can switch to the sub-organization to use the **Home** page of the sub-organization to view the usage and health of the assets of the sub-organization they switch to.



You can perform the following tasks from this page:

| Task | Description |
|------|---|
| 1 | View a summary of the status of the Secure Agent services running in your infrastructure. Click a service to view analytics for the runtime environment where the service runs. For information about monitoring runtime environments, see “Monitoring Secure Agents” on page 29 . |
| 2 | View a summary of data processing statistics for the last four weeks or the last six months for all services and domains in your infrastructure. Click a segment in the bar chart to view analytics for a service or domain type. |
| 3 | Click a location to view the status of the Secure Agents or domains within the location. The assets within the location appear in a pop-up panel. Click an asset name to view details. |
| 4 | Click the icon to add a new location to the map, or to edit an existing location. See “Adding a Secure Agent or domain to a map location” on page 28 for details. |

| Task | Description |
|------|---|
| 5 | If your organization uses other services in Informatica Intelligent Cloud Services, the page displays an overview of the analytics collected for each service. Click a value in the panel to view detailed analytics. For more information about monitoring Informatica Intelligent Cloud Services, see Chapter 4, "Monitor Informatica Intelligent Cloud Services Data Integration" on page 36 . |
| 6 | Click the notifications icon to view alerts for issues that occur with an Operational Insights resource. |

Adding a Secure Agent or domain to a map location

You can organize Secure Agents and domains by geographical location on the interactive map that is displayed on the Operational Insights **Home** page. Assigning Secure Agents and domains to locations helps you analyze performance and determine capacity and processing capabilities across the enterprise. Parent organization users can switch to a sub-organization to organize the Secure Agents by geographical location for the sub-organization.

To add a Secure Agent or domain to the map:

1. Click the edit (pencil) icon on the map.
2. Select a Secure Agent or domain from the list, and then click **Pin**.
3. Position the pin where you want to add the location on the map.
To add the Secure Agent or domain to an existing location, position the pin inside the location.
4. If you are adding a new location, enter a name for the location.

Monitoring infrastructure health

Use the **Secure Agents & Groups** panel to view the status of the assets in your Informatica infrastructure. Parent organization users can switch to a sub-organization to use the **Secure Agents & Groups** panel to view the assets of the sub-organization.

The **Secure Agents & Groups** panel displays the status of the following assets:

- Each Secure Agent running in each runtime environment.
- The services running on each Secure Agent.
- The nodes and services running in each Informatica domain.
- The Informatica Intelligent Cloud Services running in the data center, which is also called a POD (Point of Deployment), that your organization uses.

Complete the following steps to view the status of the assets in your Informatica infrastructure:

1. Click **All Infrastructure** on the Operational Insights navigation bar.
The **Secure Agents & Groups** panel displays the status of Informatica assets.
2. Click a runtime environment, a Secure Agent, or an Informatica domain to view details.

Monitoring the OI Data Collector service

You can monitor the OI Data Collector service that runs on a Secure Agent.

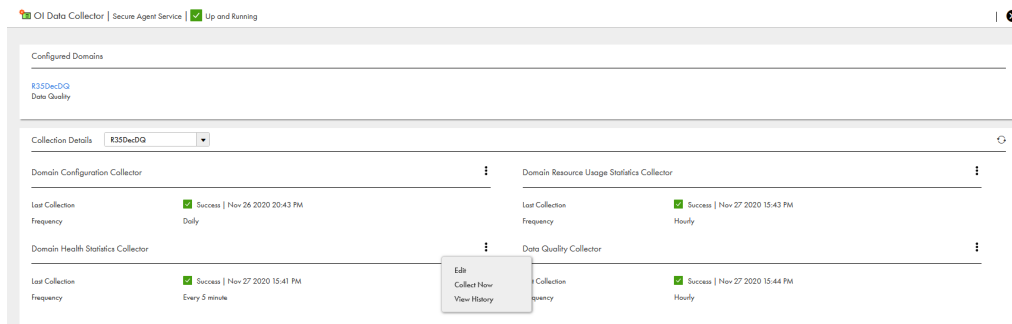
The OI Data Collector service runs the data collectors that collect the operational data and domain-related metadata used by Operational Insights. The Secure Agent uploads the collected operational data and domain-related metadata to Informatica Intelligent Cloud Services.

The page displays the following data for the OI Data Collector service running on a Secure Agent:

- The status of the service.
- The list of domains that the data collectors the OI Data Collector service runs collect data from.
- The list of data collectors enabled for each domain, and the most recent collection status.

Complete the following steps to view the status of the OI Data Collector service:

1. Click **All Infrastructure** in the navigation bar on the left side of the page.
The All Infrastructure page displays the status of the Secure Agents and Secure Agent services running in each runtime environment.
2. Click a Secure Agent that runs the OI Data Collector service.
3. Click the **OI Data Collector** link.
4. From the menu for a collector, you can select the following options:
 - Select **Edit** to modify the collector configuration.
 - Select **Collect Now** to trigger an on-demand data collection.
 - Select **View History** to view the history of the on-demand data collection.



Note: You can refresh to get the latest collection details.


Zooming in on graph details

You can zoom in on graphs to view Secure Agent resource utilization details for a specific time frame. You select the start and end times to zoom in on in the graphs.


To zoom in on graph details, complete the following steps:

1. Place the cursor on the time frame start point in a graph.
2. Left-click your mouse.
3. Drag the cursor to the time frame end point in the graph.

The following image shows the time frame from 10:51 p.m. to 12:51 p.m. selected:



The resource utilization graphs update to display data only for the specified time frame, as shown in the following image:



4. Click **Reset zoom** to return the graphs to the original state.

Infrastructure alerts

You can configure Operational Insights to send email notifications when an issue occurs within a domain or with a Secure Agent.

You can configure alerts for the following events:

- The domain or the Secure Agent is unavailable.
- A collector, service, or node running within the domain is unavailable.
- CPU or memory consumption by a domain node, a domain service, or the Secure Agent host has crossed a configurable threshold.
- Disk usage by a Secure Agent has crossed a configurable threshold.

Note: You can view the alert notifications on the **Home** page.

You can view Informatica assets for organizations on the **Infrastructure Alerts** panel of the **Alerts** page. Parent organization users with sub-organization access privileges can switch to the sub-organization to view those assets.

You can enable or disable individual alerts or all alerts for each domain or Secure Agent that Operational Insights monitors. You also specify the Informatica Intelligent Cloud Services users, user group names, or email recipients outside the organization that receive alert notifications.

You can create alert scripts that Operational Insights executes to perform additional actions when an alert is triggered. See [“Use alert scripts” on page 33](#) for more information about creating and using alert scripts.

Configuring infrastructure alerts

When you configure the infrastructure alerts, you can configure the source of the alerts: domain, Secure Agent, or Secure Agent service. You can also include recipients for the alerts and optionally include a script for additional actions.

1. Click **Alerts** in the left navigation bar.
2. Click the **Infrastructure Alerts** tab.
3. Select a domain, Secure Agent, or Secure Agent service.
4. Configure each alert that you want to enable.
5. Define the recipients that receive email notifications when an issue occurs. Perform one of the following steps:
 - To add Informatica Intelligent Cloud Services user names, under the **Email Recipients** section, click **Add Recipients > Users**. Select users from the list and click **Add**. You can also search for users.
 - To add Informatica Intelligent Cloud Services user groups, under the **Email Recipients** section, click **Add Recipients > User Groups**. Select user groups from the list and click **Add**. You can also search for user groups. When you add a user group, all the users who belong to the group receive the alert.
 - To add an external email address that belongs to a user outside the organization, under the **Email Recipients** section, click **Add Recipients > External Emails**. Click **+** and enter a fully qualified and valid email address. Click **Add**. You can also add multiple external email addresses. The email addresses must be unique.

The **Email Recipients** section displays the selected users, user groups, and external email addresses.

Note: If you switch to a sub-organization from a parent organization to configure alerts, Operational Insights doesn't show all the parent organization users that are subscribed to the sub-organization. If a user doesn't appear on the email recipients list, enter the email address manually.

6. If you want to use an alert script that Operational Insights executes to perform additional actions when an alert is triggered, select the **Run Custom Script** check box, and then enter the path to the script file on the Secure Agent host.

For more information about using alert scripts, see [“Configuring a Secure Agent to use an alert script” on page 34](#).

7. Optionally, configure alerts for Secure Agent services.

For information about configuring alerts for services, see [“Configuring alerts for Secure Agent services” on page 35](#).

Use alert scripts

You can create scripts that Operational Insights executes to perform additional tasks when an alert is triggered, such as creating a support ticket or taking a snapshot of CPU statistics.

Operational Insights contains rules that correspond to the Secure Agent alerts that you enable. You specify the name of each rule to execute in the script. For each rule, you specify parameters that Operational Insights passes to the script.

An alert script can invoke a batch file, a shell script, or an EXE file. The executable file or program must be self contained.

Copy the script to a directory on the Secure Agent host. Each time the script executes, Operational Insights writes the script output and errors to log files in the same directory on the Secure Agent host where the script executes.

If the script output directory is under the Secure Agent installation directory, you might see an error when you open the log file on Windows. To resolve this issue, perform one of the following steps:

- Open the log file with the Administrator privilege.
- Grant read, write, and execute permissions for the script output directory and the files within it.

Alternatively, you can create the script output directory outside the Secure Agent installation directory and manually move the earlier log files to the new script output directory. In Administrator, open the **Runtime Environments** page. Edit the Secure Agent and select the **OI Data Collector** service. Edit the **scriptLogDir** property value for OpsInsights and specify a script output directory that is outside the Secure Agent installation directory.

[Click here](#) to view an example of an alert script.

Alert script rules and parameters

Add the rule name and the rule parameters that for the rules you want to specify in the script file. Specify parameters as key value pairs.

You can configure multiple rules in a script.

The following table describes the rule names and parameters associated with Secure Agent alerts.

| Secure Agent Alert | Rule Name | Rule Parameters |
|---|-----------------------------------|---|
| Secure Agent is unavailable for a duration of 15 minutes. | secure-agent-unavailable-rule | "ruleName": "secure-agent-unavailable-rule", "timestamp": "\${timestamp}", "agentName": "\${agentName}" |
| CPU usage by the Secure Agent exceeds xx % for a duration of 30 minutes. | secure-agent-cpu-overused-rule | "ruleName": "secure-agent-cpu-overused-rule", "timestamp": "\${timestamp}", "agentName": "\${agentName}", "actualUsage": "\${actualUsage}", "thresholdValue": "\${thresholdValue}" |
| Memory usage by the Secure Agent exceeds xx % for a duration of 30 minutes. | secure-agent-memory-overused-rule | "ruleName": "secure-agent-memory-overused-rule", "timestamp": "\${timestamp}", "agentName": "\${agentName}", "actualUsage": "\${actualUsage}", "thresholdValue": "\${thresholdValue}" |

Configuring a Secure Agent to use an alert script

Complete the following steps to configure a Secure Agent to use an alert script.

1. Copy the script file to the Secure Agent host.
2. Click **Alerts** in the left navigation bar.
3. Click **Infrastructure Alerts**.
4. Select the Secure Agent that uses the alert script.
5. Select the **Run Custom Script** check box.
6. Click the **Edit** link, and then enter the path to the script file on the Secure Agent host.
7. Click the **Save** link.

Purging Secure Agent script logs

You can specify the number of Secure Agent script log files to retain in the directory where the script runs. The application retains the most recent log files up to the value specified, and deletes all older log files.

By default, the application retains the 50 most recent script log files.

1. Click **Secure Agents**.
2. Select the Secure Agent where the Ops Insights Collector runs.
3. Enter the number of log files to keep as the value for the scriptLogRetentionCount property.

Configuring alerts for Secure Agent services

You can configure Operational Insights to send alerts when an issue occurs with individual services on a Secure Agent.

You can configure alerts for the following events for a Secure Agent service:

- The service is unavailable
- CPU usage by the service crosses a configurable threshold
- Memory usage by the service crosses a configurable threshold

To configure alerts for Secure Agent services, perform the following steps:

1. On the **Infrastructure Alerts** tab of the **Alerts** page, select the Secure Agent.
2. If alerts are disabled for the Secure Agent, enable alerts.
 - a. Click **Enable All**.
 - b. Enable or disable individual Secure Agent alerts.
3. Click **Add Service Specific Alert**.
4. Select the Secure Agent service that you want to configure alerts for.
5. Configure each alert that you want to enable.
6. Enter an Informatica Intelligent Cloud Services user, user group names, or email recipients outside the organization that receive email notifications when an issue occurs.
7. If you want Operational Insights to perform additional actions when an alert is triggered, enable **Run Custom Script**, and then enter the path to the script file on the Secure Agent machine.
8. Repeat steps [3](#) through [7](#) for any additional Secure Agent services.

CHAPTER 4

Monitor Informatica Intelligent Cloud Services Data Integration

If an organization uses Informatica Intelligent Cloud Services Data Integration, you can use Operational Insights to view analytics for Data Integration assets, Data Integration jobs, details of specific jobs, job history of specific assets, scheduled jobs, and Data Integration connections. View analytics for Data Integration on the **Data Integration** page.

You can monitor the following types of Data Integration assets:

- Command tasks
- Data transfer tasks
- Dynamic mapping tasks
- Mappings and mapping tasks
- Masking tasks
- PowerCenter tasks
- Replication tasks
- Synchronization tasks
- Taskflows and linear taskflows


For more information about using Data Integration, see the Data Integration help.

To monitor Data Integration assets with Operational Insights, you need the appropriate license. For more information or to request this feature, contact Informatica Global Customer Support.

View Data Integration analytics

Use the **Overview** tab of the **Data Integration** page to show analytics for the jobs that have run in the organization in the last 24 hours. The panels on the **Overview** tab provide insights into specific areas of data processing and job run analytics.

The following image shows the **Overview** tab:



Hover over a bar in any chart on the **Overview** tab to view summary details. To view detailed information about the jobs completed in the selected time period, click the bar on the graph.

Information about jobs is current as of the most recent data refresh. Data is refreshed approximately every hour at the top of the hour.

The **Overview** tab contains the following panels:

Jobs

Shows an overview of analytics for the jobs that are running or have run in the organization in the last 24 hours. The **Historical** panel shows completed jobs. Click **Recent** to view currently running jobs and recently completed jobs. You can refresh the **Recent** panel manually.

Click a number to view the relevant jobs on either panel. The figures shown are for the 24 hours prior to the most recent data refresh.

Activity

Shows the number of rows processed and the number of jobs run in the selected time period. You can choose to view data for the past 24 hours or for the past 7 days.

Click **Detailed Charts** to show analytics on the total rows of data processed and number of jobs run during the selected time period.

Top Secure Agents

Shows the number of jobs run for the most used Secure Agents in the organization. Hover over the graph to view the jobs according to status. You can choose to view data for the past 24 hours or for the past 7 days.

Click **Detailed Charts** to show the rows of data processed and the number of jobs run by the selected Secure Agents during the selected time period.

Top Targets and Top Sources

The **Top Targets** panel shows a summary of data rows written by the five most used connections in the organization. To view the **Top Sources** panel, select **Sources** from the menu. The **Top Sources** panel shows a summary of data rows read by the five most used connections in your organization. You can choose to view data for the past 24 hours or for the past seven days.

Click **Detailed Charts** in either the **Top Targets** or **Top Sources** panel to show analytics on the total rows of data written or read, and the number of jobs run for the selected connectors during the selected time period.

Top Projects and Top Folders

The **Top Projects** panel shows a summary of jobs run for the top five projects in the organization. Hover over the graph to view the jobs according to status. To view the **Top Folders** panel, select **Folders** from the menu. You can choose to view data for the past 24 hours or for the past seven days.

Click **Detailed Charts** on either the **Top Projects** panel or the **Top Folders** panel to show analytics on the total number of rows read and jobs run for the selected objects during the selected period of time.

Data Volume

Shows a weekly summary of rows processed every hour. You can view data for the following time periods:

- 1 week actual
- 2 week average
- 1 month average

Hover over the chart to view job information for a day and time. If you view actual data for one week, you can view details about the jobs that completed in that time period. Click the rectangle to view details about the jobs.

When you view actual data for one week, you can toggle to view actual data for the past month. Default is the first week of the month.

View Data Integration jobs

Click the **Jobs** tab of the **Data Integration** page to view details about the jobs that have run in the organization. The **Jobs** tab is divided into a **Historical** panel to view completed jobs and a **Recent** panel to view currently running jobs and recently completed jobs. Click **Download log** to download the session logs of jobs.

Historical panel

The following image shows the **Historical** panel:

The screenshot shows the 'Data Integration' interface with the 'Jobs' tab selected. The 'Historical' panel is active, displaying a table of jobs. The table has columns for Job, Asset, Location, Project Name, Secure Agent, Start Time, End Time, Rows Processed, Started By, Parent Asset Name, and Status. The table is filtered to show jobs from the last 24 hours. A 'Download Log' button is visible next to the first job row. Numbered callouts 1 through 8 point to various UI elements: 1 points to the filter dropdown, 2 points to the 'Add Field' button, 3 points to the sort dropdown, 4 points to the filter icon, 5 points to the settings icon, 6 points to the 'Export' button, 7 points to the 'Download Log' button, and 8 points to the alert icon.

| Job | Asset | Location | Project Name | Secure Agent | Start Time | End Time | Rows Processed | Started By | Parent Asset Na... | Status |
|----------------|--------------------|----------------|----------------|----------------|------------------|------------------|----------------|-------------|--------------------|---------|
| massingest... | massingestion_F... | _CDIOL_New1... | _CDIOL_New1... | n/a | May 31 2021 0... | May 31 2021 0... | 2 | netta_cdioi | | Success |
| PowerCente... | PowerCenterTas... | _CDIOL_New1... | _CDIOL_New1... | ILW1PF18W2P... | May 31 2021 0... | May 31 2021 0... | 4 | 1 hour | | Success |
| Synchroniza... | Synchronization... | _CDIOL_New1... | _CDIOL_New1... | ILW1PF18W2P... | May 31 2021 0... | May 31 2021 0... | 4 | 1 hour | | Success |
| echo-2351... | echo | _CDIOL_New1... | _CDIOL_New1... | n/a | May 31 2021 0... | May 31 2021 0... | 0 | netta_cdioi | commandTask... | Success |
| MappingTas... | MappingTaskFF... | _CDIOL_New\... | _CDIOL_New | ILW1PF18W2P... | May 31 2021 0... | May 31 2021 0... | 1 | 1 hour | | Success |
| massingest... | massingestion_F... | _CDIOL_New1... | _CDIOL_New1... | n/a | May 31 2021 0... | May 31 2021 0... | 2 | netta_cdioi | | Success |
| MappingTas... | MappingTaskFF... | _CDIOL_New\... | _CDIOL_New | ILW1PF18W2P... | May 31 2021 0... | May 31 2021 0... | 1 | test | | Success |
| MappingTas... | MappingTaskFF... | _CDIOL_New1... | _CDIOL_New1... | n/a | May 31 2021 0... | May 31 2021 0... | 0 | 1 hour | | Failed |
| massingest... | massingestion_F... | _CDIOL_New1... | _CDIOL_New1... | n/a | May 31 2021 0... | May 31 2021 0... | 2 | SCHEDULER | | Success |

1. Change the time period. You can view jobs run today, the last 24 hours, the last week, or enter a custom range.
2. Add a new filter.
3. Sort the jobs on the page. You can sort by start time or end time.
4. Add or remove filters.
5. Open the Settings window.
6. Open the Exports panel. You can export up to 10,000 rows of jobs data.
7. Download session log.
8. Create alert.

If a task contains child jobs, Operational Insights lists the parent and child jobs as separate jobs. For example, a dynamic mapping task contains three jobs. Operational Insights lists the dynamic mapping task and the three jobs on the **Jobs** tab. To view information about a child job, click the job name on the **Jobs** tab. When you view parent job details from the **Historical** panel, you cannot view information about a child job in the parent job details.

The **Historical** panel displays child jobs for taskflows and linear taskflows but does not show parent jobs.

By default, the **Historical** panel shows jobs that completed in the last 24 hours. You can edit the **End Time** filter to show jobs that were completed from the last hour to the last 30 days. You can apply the following filters:

- Asset
- Asset Type
- End Time

- Location
- Parent Asset Name
- Project Name
- Runtime Environment
- Started By
- Status
- Secure Agent

By default, the following properties show for each job:

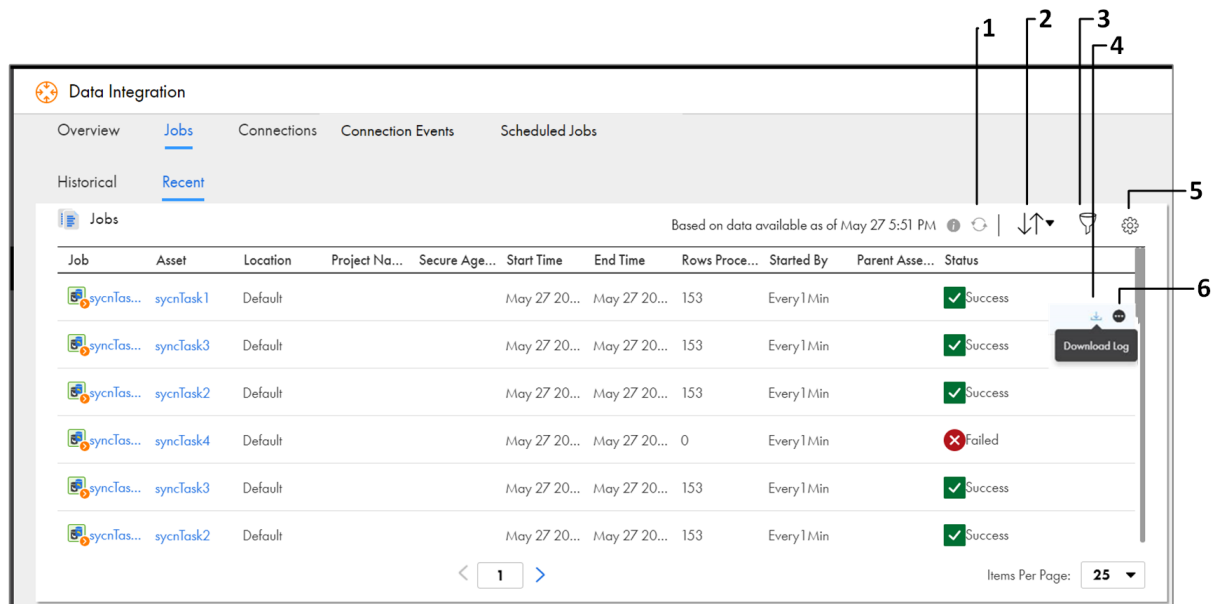
- Job
- Asset
- Location
- Project Name
- Secure Agent
- Start Time
- End Time
- Rows Processed
- Started By
- Parent Asset Name
- Status

You can show additional properties by right-clicking on the column heading area and selecting any of the following filters:

- Asset Type
- Subtasks
- Runtime Environment
- Duration
- Success Rows
- Failure Rows
- Started By
- Error Message

Recent panel

The following image shows the **Recent** panel:



1. Refresh jobs to the current status.
2. Sort the jobs on the page. You can sort by start time or end time.
3. Add or remove filters.
4. Download session log.
5. Open the Settings window.
6. Create alert.

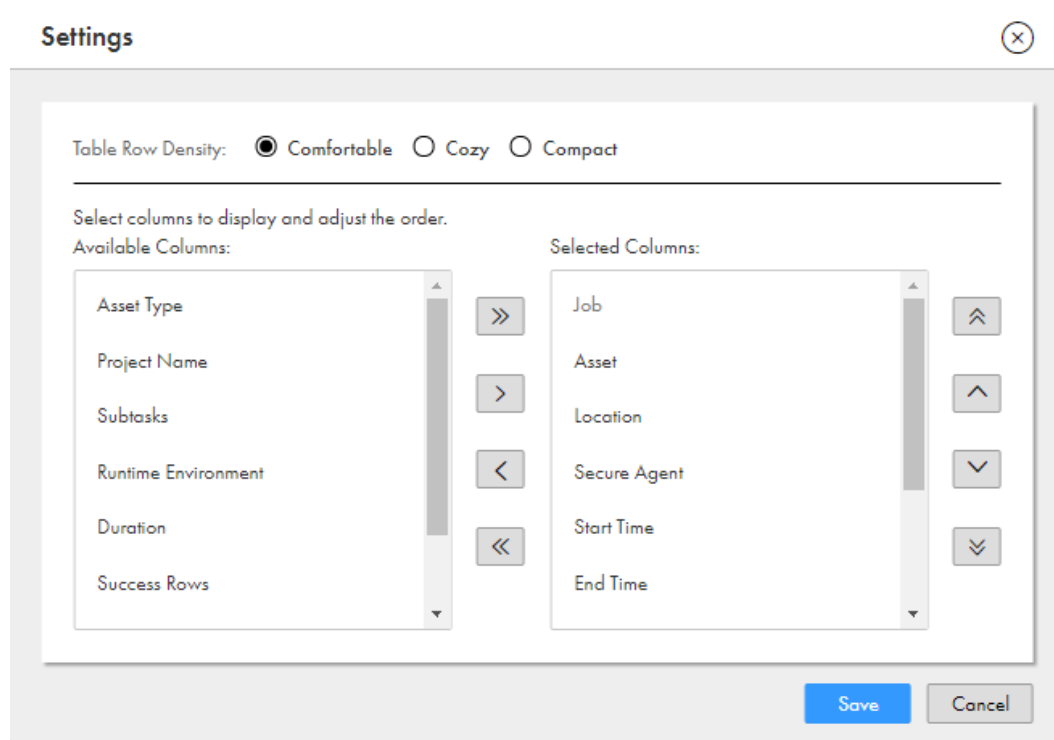
If a task or taskflow contains child jobs, Operational Insights lists the parent and child jobs as separate jobs. For example, a dynamic mapping task contains three jobs. Operational Insights lists the dynamic mapping task and the three jobs on the Jobs tab. To view information about a child job, click the job name on the **Jobs** tab or on the parent job details page.

Configure Jobs page settings

Configure the **Jobs** page properties and layout in the **Settings** window. Settings remain in place the next time you log in.

To open the **Settings** window, click **Settings** on the **Jobs** page.

The following image shows the **Settings** window:



The following table describes the settings that you can adjust:

| Setting | Description |
|------------------------|---|
| Jobs table row density | Adjusts the height of each row and space between entries. You can choose the following densities: <ul style="list-style-type: none"> - Comfortable. Maximum row height. - Cozy. Medium row height. - Compact. Minimum row height. |
| Available columns | Determines the columns to display on the Jobs page. Select a column name and use the left/right arrows to move it in or out of the Selected Columns area. |
| Column order | Arrange the selected columns in the order that you want them to appear on the Jobs page. Select a column name and use the up/down arrows to adjust the column position. |

After you have configured the page settings, click **Save**.

Exporting jobs data

Export jobs data from the **Historical** panel on the **Jobs** page to a CSV file. After you export the data, you can download the file or email it as an attachment.

When you export jobs data, Operational Insights exports the data with the current page filters applied. Operational Insights exports job start and end times as strings in Coordinated Universal Time.

1. Click **Export**.
2. In the **Export Job Data** window, enter a file name.

3. Click **Export**.
4. To view details about the export job, click the file name in the **Exports** panel.
The **Export Job Details** page displays information about the export job properties and filter criteria.
5. To download the file, perform one of the following actions:
 - In the **Exports** panel, click **Download**.
 - On the **Export Job Details** page, click **Download**.
6. To email the file, in the **Exports** panel, click **Email**.
Operational Insights sends the file to the email address associated with your Informatica Intelligent Cloud Services account.
If the email does not appear in your inbox, check your spam folder.

View details for a specific job

You can drill down on a specific job to view details about the job and to download a session log.

To view the job details, click the job name on the **Job** column.

The following image shows the job details for a mapping task:

The screenshot displays the job details for a mapping task. The page is titled "MappingTaskFF_FF-1159". It is divided into two main sections: "Job Properties" and "Results".

Job Properties:

- Asset: MappingTaskFF_FF
- Instance ID: 1159
- Asset Type: Mapping Task
- Started By: test
- Start Time: May 31, 2021, 12:30:21 PM
- End Time: May 31, 2021, 12:30:36 PM
- Duration: 15 seconds
- Runtime Environment: ILW1PF18W2PM-1
- Secure Agent: ILW1PF18W2PM_changed

Results:

- Status: Success
- Success Rows: 1
- Error Rows: 0
- Error Message:
- Session Log: [Download Session Log](#)

Individual Source/Target Results:

| Name | Success Rows | Error Rows | Error Message |
|--------------------|--------------|------------|---------------|
| Source | 1 | 0 | |
| tgt_cdiol_test_txt | 1 | 0 | |

At the bottom of the table, there is a pagination control showing "1" and an "Items Per Page" dropdown set to "25".

Click **Download Session Log** in the **Results** area to download the session log of the job.

If the job was started by a taskflow, dynamic mapping task, or replication task, you can view details about the taskflow or task. Click the taskflow or task name in the **Job Properties** area.

The following image shows the details for a replication task:

The screenshot displays the 'Account-8' job details page. It is divided into two main sections: 'Job Properties' and 'Results'.

Job Properties:

- Asset: Account
- Instance ID: 8
- Asset Type: Replication Task
- Stop on Error: Cancel processing the remaining objects
- Started By: ReplicationTask_salesforceWithEmailNotification-8
- Start Time: May 31, 2021, 02:10:46 PM
- End Time: May 31, 2021, 02:11:13 PM
- Duration: 27 seconds
- Runtime Environment: ILW1PF18W2PM-t1
- Secure Agent: ILW1PF18W2PM_changed

Results:

- Status: ✔ Success
- Success Rows: 734
- Error Rows: 0
- Error Message: No errors encountered.
- Session Log: [Download Session Log](#)

Individual Source/Target Results:

| Name | Success Rows | Error Rows | Error Message |
|------------------|--------------|------------|---------------|
| SQL_Account | 734 | 0 | |
| SF52_ACCOUNT_csv | 734 | 0 | |

At the bottom right, there is a pagination control showing 'Items Per Page: 25'.

If a job fails, the **Insights** area provides links to documentation to help you resolve the error.

The following image shows the Insights area:

▼ Insights(5)

ERROR: "FR_3000 Error opening file [C:\Twitter Connection\Src_Twitter.csv]. Operating system error message [The system cannot find the path specified.], every 2nd run"

... fails with the error FR_3000 Error opening file [C:\Twitter Connection\Src_Twitter. ... Operating system error message [The system cannot find the path specified.], every second run and the ...

[View Article on Informatica Network](#)

ERRRR: "FR_3000 Error opening file [.\dummy_data.csv]. Operating system error message [The system cannot find the file specified.]" when running a task which has saved Query as source in Informatica Cloud

... dummy_data_csv' and can be noticed in the session log instead of the actual query Secure agent is unable to access SQ_dummy_data_csv file or the file is not available in the desired location.

[View Article on Informatica Network](#)

"FR_3000 Error opening file [<path of lookup file>]. Operating system error message [The system cannot find the path specified.]" for an DSS task with mapplet having Lookup transformation in Informatica Cloud

Operating system error message [The system cannot find the path specified.] This errors occurs when the mapplet used by the DSS task contains a Lookup transformation, and the lookup file does ...

[View Article on Informatica Network](#)

"FR_3000 Error opening file [filepath]. Operating system error message [The system cannot find the path specified.]" while using a Flat File connection for a DSS task in Informatica Cloud

... Service (DSS) task, the following error message is displayed: FR_3000 Error opening file [filepath]. ... Operating system error message [The system cannot find the path specified.].

[View Article on Informatica Network](#)


To view error remediation recommendations, you need the appropriate license.

View job history for a specific asset

You can view the job history for a specific mapping or task. Use the **Job History** page to view job run analytics such as run time and completion status.

To view the job history for an asset, click the asset name on the **Jobs** tab. The **Job History** page displays details about each instance of the mapping or task.

The following image shows the **Job History** page:



1. Filter the page by end time or status. You can view data for the past 24 hours or the last seven days. You can apply the following additional filters:

- Runtime Environment
- Started By

2. To view jobs with a particular status, click the status.

3. To view row information for a job, hover over a point on the graph.

4. Job instances. You can sort the jobs list by start time or end time.

By default the **Jobs** area displays the following properties for each job:

- Job
- Started By
- Parent Asset Name
- Start Time
- End Time
- Rows Processed
- Status

You can display the following additional properties by right clicking on the column heading area:

- Location
- Subtasks
- Runtime Environment
- Duration
- Success Rows
- Failure Rows
- Error Message

5. To view the **Job Details** page for a job, click the job name in the **Jobs** area.

View Data Integration connections

Click the **Connections** tab to view analytics for the connections in the organization.

Each time a connection is accessed, Operational Insights logs the connection details on the **Connections** tab. The information on the tab is updated every hour.

Operational Insights groups the connections in your organization based on the connector type. In the **Connector Type** area, Operational Insights shows the number of active and the number of idle connections of each type. An active connection was used to run at least one job in the last 33 days. An idle connection was not used to run a job in the last 33 days.

By default, Operational Insights displays details for the most common connector type in your organization. To view details about the connections of a specific type, click the connector type in the **Connector Type** area. The **No. of Connections** area displays the number of active and idle connections for the connector type. Click a number to view aggregated details of the active connections in the **Connections** table from the last 24 hours. Click the connector type again to deselect it and view details for all connections accessed in the last 24 hours.

Click the number of active connections to view connection usage trends on the **Connection Usage Trend** area. You can select up to five connections from the **Connections** table to view. View connection usage data on the **Rows Read** and the **Rows Written** charts according to time duration from one day through one month.

You can view up to 16 connector types at a time. To select the connector types to view, click **View More**.

You can apply filters to see connections in a specific runtime environment or connections that use a specific Secure Agent.

The following image shows the **Connections** tab filtered to show flat file connections:

The screenshot displays the Informatica Operational Insights interface. The left sidebar shows navigation options: Home, Alerts, All Infrastructure, Data Profiling, Mass Ingestion, and Data Integration (selected). The main content area is titled 'Data Integration' and has tabs for Overview, Jobs, Connections (selected), Connection Events, and Scheduled Jobs. Under the 'Connections' tab, there's a 'Connector Type' section with filters for Flat File (31/204), Oracle (3/3), Salesforce (2/2), Amazon S3 v2 (0/1), SAP_ALE_IDoc_VW... (0/1), SFTP (0/1), and SqlServer2014 (1/1). Below this, it shows 'No. of Connections' with 31 Active and 173 Idle. A 'Connection Usage Trend' section is partially visible. The main table shows a list of connections, filtered to Flat File connectors. The table has columns: Connection, Connector, Runtime Environment, Secure Agent, Last Accessed, Rows Read, Rows Written, and Status. The first row is selected, showing a connection named 'FF_test' with 241 rows read and 241 rows written.

| Connection | Connector | Runtime Environment | Secure Agent | Last Accessed | Rows Read | Rows Written | Status |
|------------|-----------|--------------------------------|--------------------------------|----------------------|-----------|--------------|--------|
| FF_test | Flat File | tswh7cdsioagn1.informatica.com | tswh7cdsioagn1.informatica.com | Nov 03 2022 11:55 AM | 241 | 241 | Active |
| FF_test | Flat File | tswh7cdsioagn1.informatica.com | tswh7cdsioagn1.informatica.com | Nov 03 2022 11:55 AM | 11 | 11 | Active |
| FF_test | Flat File | tswh7cdsioagn1.informatica.com | tswh7cdsioagn1.informatica.com | Nov 03 2022 10:30 AM | 10 | 10 | Active |
| FF_test | Flat File | tswh7cdsioagn1.informatica.com | tswh7cdsioagn1.informatica.com | Nov 03 2022 09:30 AM | 10 | 10 | Active |
| FF_test | Flat File | tswh7cdsioagn1.informatica.com | tswh7cdsioagn1.informatica.com | Nov 03 2022 08:30 AM | 10 | 10 | Active |
| FF_test | Flat File | tswh7cdsioagn1.informatica.com | tswh7cdsioagn1.informatica.com | Nov 03 2022 07:30 AM | 10 | 10 | Active |
| FF_test | Flat File | tswh7cdsioagn1.informatica.com | tswh7cdsioagn1.informatica.com | Nov 03 2022 06:30 AM | 10 | 10 | Active |
| FF_test | Flat File | tswh7cdsioagn1.informatica.com | tswh7cdsioagn1.informatica.com | Nov 03 2022 05:30 AM | 10 | 10 | Active |
| FF_test | Flat File | tswh7cdsioagn1.informatica.com | tswh7cdsioagn1.informatica.com | Nov 03 2022 04:30 AM | 10 | 10 | Active |

View Data Integration connection events

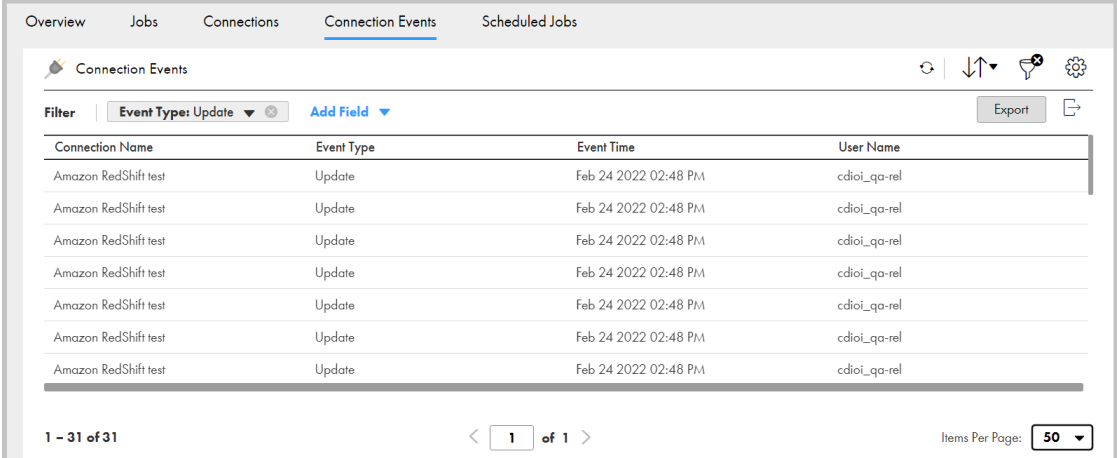
Click the **Connection Events** tab to download and view reports of connection events in the organization.

Each time Operational Insights triggers a connection event, Operational Insights logs the event details on the **Connection Events** tab.

By default, the **Connection Events** tab lists Create, Update, and Delete connections according to the time they are accessed. You can filter the view to show events for a selected event type, connection name, or user name.

You can export reports of the connection events according to the filters you apply. You can export up to 10,000 events.

The following image shows the **Connection Events** tab filtered to show Update events:




The screenshot shows the 'Connection Events' tab in a web application. The navigation bar includes 'Overview', 'Jobs', 'Connections', 'Connection Events' (selected), and 'Scheduled Jobs'. The main content area has a 'Filter' section with 'Event Type: Update' selected and an 'Add Field' button. An 'Export' button is also visible. The table below lists 8 events, all of type 'Update' for 'Amazon RedShift test' connections, occurring on 'Feb 24 2022 02:48 PM' by user 'cdioi_qa-rel'. The footer shows '1 - 31 of 31' items, page '1 of 1', and 'Items Per Page: 50'.

| Connection Name | Event Type | Event Time | User Name |
|----------------------|------------|----------------------|--------------|
| Amazon RedShift test | Update | Feb 24 2022 02:48 PM | cdioi_qa-rel |
| Amazon RedShift test | Update | Feb 24 2022 02:48 PM | cdioi_qa-rel |
| Amazon RedShift test | Update | Feb 24 2022 02:48 PM | cdioi_qa-rel |
| Amazon RedShift test | Update | Feb 24 2022 02:48 PM | cdioi_qa-rel |
| Amazon RedShift test | Update | Feb 24 2022 02:48 PM | cdioi_qa-rel |
| Amazon RedShift test | Update | Feb 24 2022 02:48 PM | cdioi_qa-rel |
| Amazon RedShift test | Update | Feb 24 2022 02:48 PM | cdioi_qa-rel |
| Amazon RedShift test | Update | Feb 24 2022 02:48 PM | cdioi_qa-rel |

View scheduled jobs

Click the **Scheduled Jobs** tab of the **Data Integration** page to view the scheduled jobs in the organization. The **Scheduled Jobs** tab shows a graph of the number of scheduled jobs per hour for each runtime environment for the selected time frame. It also shows the job details for each job in the selected time frame.

The following image shows the **Scheduled Jobs** tab:



1. Filter the page. You can filter by time frame, runtime environment, or schedule.
2. Hover over the graph to view detailed information about the scheduled jobs.
3. Click a runtime environment to hide or show it on the graph.
4. Export a list of scheduled jobs with the current filters applied.
5. Open the Exports panel.
6. Sort the jobs table by scheduled time.
7. Adjust the table row density.

To view details about a job, click the job name.

You can export a list of your scheduled jobs as a CSV file with the current page filters applied. To export your scheduled jobs, click **Export**. After the export completes, you can download the file by opening the **Exports** panel.

Data Integration alerts

You can configure Operational Insights to send alert notifications for Data Integration jobs. Configure alerts for specific jobs in the organization.

For example, you might want to configure an alert to notify you when the mapping tasks in a project have been running for more than five minutes. Or, you might want to receive an alert when a task that loads order data to a target fails.

To configure alerts, you must have the Admin or Operator user role, or a custom user role with the Data Integration Job Alerts privilege for Operational Insights.

You can configure alerts for the following Data Integration assets:

- Mapping tasks

- Synchronization tasks
- Dynamic mapping tasks
- Data transfer tasks
- Replication tasks
- PowerCenter tasks
- Linear taskflows

You can configure alerts for the following events:

- The job is in a specified state.
- The duration of the job has crossed a configurable threshold.
- The number of rows processed has crossed a configurable threshold.
- The error row count has crossed a configurable threshold.

You can also configure Operational Insights to take the following actions when it sends an alert:

- Email specified users or user groups.
- Restart failed jobs.
- Stop running jobs.

Operational Insights polls the Informatica Intelligent Cloud Services repository for job data every three minutes. Operational Insights doesn't send alerts for running jobs that run for less than three minutes if the job has already completed when Operational Insights gets the job data.

Operational Insights sends email alerts for the first 550 alerts per hour, 4000 alerts per day, and 28000 alerts per week in your organization. To change the maximum number of alerts that Operational Insights sends, contact Informatica Global Customer Support.

Configuring alerts for Data Integration jobs

Configure alerts for Data Integration jobs on the **Data Integration Alerts** tab.

1. Click **Create Alert**.
2. Configure the alert details:
 - a. If you do not want to enable the alert right away, disable the alert.
 - b. Enter a name for the alert. Optionally, enter a description of the alert.
 - c. Configure the alert scope. Perform one of the following actions:
 - To apply the alert to a task, select **Task Asset** and click **Select...**. Select an asset and click **Select**.
 - To apply the alert to a project or a folder, select **Project or Folder** and click **Select...**. Select a project or folder and click **Select**.
3. Configure the alert condition.
4. Enter an Informatica Intelligent Cloud Services user or user group names that receive email notifications for the alert.

To receive alert emails from a sub-organization that you subscribe to, create a user in the sub-organization. Configure Operational Insights to send alert emails to the user in the sub-organization.
5. Configure other alert actions based on the alert condition.
6. Click **Save**.

CHAPTER 5

Monitor Informatica Intelligent Cloud Services Application Integration

If your organization uses Application Integration, you can use Operational Insights to view analytics for Application Integration assets. To monitor Application Integration assets with Operational Insights, your organization must have the Application Integration license and be enabled to view Application Integration metrics.

Operational Insights offers several charts that help you quickly and visually assess the status and usage of your Application Integration assets and take appropriate corrective actions as necessary. You can view analytics related to API calls, processes, connections, and licensing. The data in the charts refreshes every 10 minutes. For more information about using Application Integration, see the Application Integration help.

To view analytics for Application Integration, open the Operational Insights service and click **Application Integration** on the Operational Insights navigation bar.

Note: Informatica is introducing Operational Insights support for Application Integration in phases. Operational Insights will be enabled first on a regional basis and then to specific customers upon request before enabling it broadly for all Application Integration customers. If you would like to use Operational Insights with Application Integration before it is broadly released, contact your Customer Success Manager.

View overall usage and health of Application Integration assets

Click the **Overview** tab of the **Application Integration** page to view overall analytics and assess the usage and health of your Application Integration assets.

You can gain insights into the following metrics:

- Number of completed and faulted processes
- Number of incoming API calls
- Overall API response time
- Processes whose API response time exceeds the average API response time
- Top 10 running processes in your organization

Note: The **Overview** tab is divided into process panels that show metric charts. The charts exclude processes that were suspended at any point of time in their lifecycle.


You can use the following panels to view the overall usage and health of Application Integration assets:

Activity

The **Activity** panel shows the total number of completed and faulted processes in your organization in the last 24 hours. The panel also shows the percentage increase or decrease in the number of completed and faulted processes since the last period.

You can use the **Activity** panel to assess whether the fault trend is moving upward or downward.

The following image shows the **Activity** panel:



Process Run by Environment

The **Process Run by Environment** panel shows the total number of completed and faulted process runs in your organization for all runtime environments. By default, the panel shows information for the last 24 hours.

You can use the **Process Run by Environment** panel to assess whether there is a need to perform load balancing across runtime environments.

The following image shows the **Process Run by Environment** panel:



1. From the list, select the time period for which you want to view the number of process runs. You can select one of the following values:

- Last 24 hours
- Last 7 days
- Last 30 days
- Custom. You can select a duration of 1 month in the past 13 months.

Hover over a bar in the chart to view the number of process runs for all runtime environments for the selected time period. Completed process runs appear in green and faulted process runs appear in red.

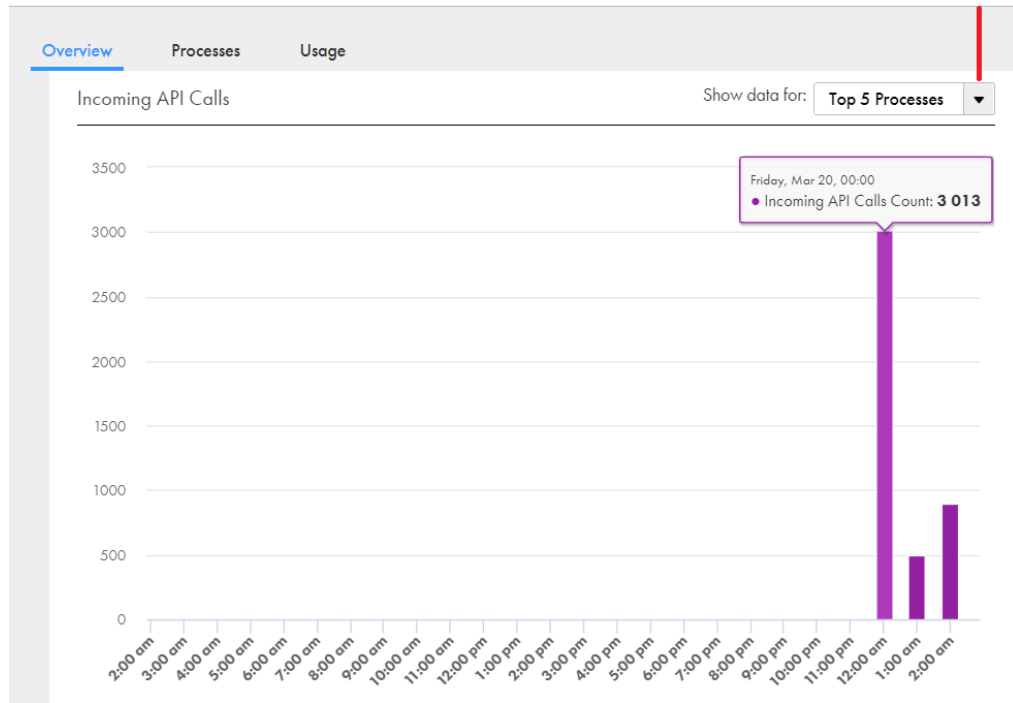
Incoming API Calls

The **Incoming API Calls** panel shows the number of hourly incoming API calls for the last 24 hours. The count also includes the incoming API calls from scheduled processes and event-based processes. By default, the panel shows information for the top 5 processes with the most incoming API calls.

You can use the **Incoming API Calls** panel for the following tasks:

- Gauge incoming traffic and identify peak and off-peak hours.
- Identify the optimal time to perform system maintenance.

The following image shows the **Incoming API Calls** panel:



1. From the list, select the process or process category for which you want to view the hourly count of incoming API calls. You can select one of the following values:

- All Processes
- Top 5 Processes
- Bottom 5 Processes
- A specific process


Hover over a bar in the chart to view the hourly count of incoming API calls.

Overall API Response Time

The **Overall API Response Time** panel shows the overall API response time in milliseconds for all completed and faulted processes, and for all runtime environments in your organization.

An increasing overall API response time indicates degraded performance probably because these APIs use complex integrations or have degraded third-party connectivity.

The following image shows the **Overall API Response Time** panel:




Hover over a point in the trend line of the chart to view the overall API response time in milliseconds.

Response Time Outliers

The **Response Time Outliers** panel shows the top 5 completed and faulted processes in your organization whose API response time is higher than the average API response time in the selected time period. By default, the panel shows information for the last 7 days.

Processes whose API response time is higher than the average API response time could have complex integrations. A comparison of the values against the average API response time in the last period helps you assess how the processes have been performing since the last time period.

The following image shows the **Response Time Outliers** panel:



1. From the list, select the time period for which you want to view the response outliers. You can select one of the following values:

- Last 24 hours
- Last 7 days
- Last 30 days
- Custom. You can select a duration of 1 month in the past 13 months.


Hover over a bar in the chart to view the API response time for the process and the average API response time for the selected time period.

Top 10 Running Processes

The **Top 10 Running Processes** panel shows the top 10 completed and faulted processes in your organization with the most runs. By default, the panel shows information for the last 24 hours.

You can use the **Top 10 Running Processes** panel to identify important processes in the organization that might need better optimization than other processes. The panel also shows the average number of process runs in the last period compared to the current period, which gives a trend indication.

The following image shows the **Top 10 Running Processes** panel:



1. From the list, select the time period for which you want to view the top 10 running processes . You can select one of the following values:

- Last 24 hours
- Last 7 days
- Last 30 days
- Custom. You can select a duration of 1 month in the past 13 months.

Hover over a bar in the chart to view the number of process runs in the current period and the average number of process runs in the last period. Use the following guidelines to understand the color coding of the bar chart:


- A grey bar indicates the average number of process runs in the last period.
- A pink bar indicates that the number of current process runs exceeds the last period's average.
- A blue bar indicates that the number of current process runs falls below the last period's average.

Note: If fewer than 10 processes were run in your organization in the selected time period, you will see fewer than 10 bars in the chart.

View Application Integration process reports

Click the **Processes** tab of the **Application Integration** page to view reports about the status and details of processes that were run in your organization.

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



By default, the **Processes** page shows all the processes that were run in your organization in the last 24 hours and across all runtime environments.

1. From the list, select a time period for which you want to view the process report. You can select one of the following values:

- Last 24 hours
- Last 7 days
- Last 30 days
- Custom. You can select a duration of 1 month in the past 13 months.

2. From the list, select a runtime environment for which you want to view the process report. You can select one of the following values:

- All Environments
- A specific runtime environment
- Cloud Only
- All Agents

3. From the left pane, select a category for which you want to view the process report. You can view all processes, completed processes, or faulted processes.

The following details are displayed for each process:

- Process name
- Number of executions
- Runtime environment in which the process was run
- Average execution time of the process in milliseconds

Note: The runtime environment displays the Secure Agent name prefixed with the Secure Agent ID in the following format:

<Secure Agent ID>_<Secure Agent Name>

Monitor usage of Application Integration assets

Click the **Usage** tab of the **Application Integration** page to monitor the usage of your Application Integration assets. You can view analytics related to API calls, processes, connections, and licensing.

View incoming API calls

Click the **APIs** tab under the **Usage** tab to view analytics related to API calls. You can view the number of daily and cumulative incoming API calls, and also analyze the number of incoming API calls based on processes and runtime environments. The count also includes the incoming API calls from scheduled processes and event-based processes.

The panels on the **APIs** tab show the API calls from parent processes only and do not include the API calls from subprocesses. The panels exclude processes that were suspended at any point of time in their lifecycle.


You can use the following panels to view incoming API call analytics:

Daily and Cumulative Incoming API Calls

The **Daily and Cumulative Incoming API Calls** panel shows the number of daily and cumulative incoming API calls. By default, the panel shows information for the current month and all runtime environments.

You can use the **Daily and Cumulative Incoming API Calls** panel to assess whether there is a surge or dip in the incoming API calls and accordingly plan for peak traffic.

The following image shows the **Daily and Cumulative Incoming API Calls** panel:



1. From the list, select the month for which you want to view the daily and cumulative incoming API calls. You can select the current month, one of the preceding 2 months, or a custom date range. If you select the **Custom** option, you can select a duration of 1 month in the past 13 months.

2. From the list, select the runtime environment or runtime environment category for which you want to view the daily and cumulative incoming API calls. You can select one of the following values:

- All Environments
- A specific runtime environment
- Cloud Only

- All Agents

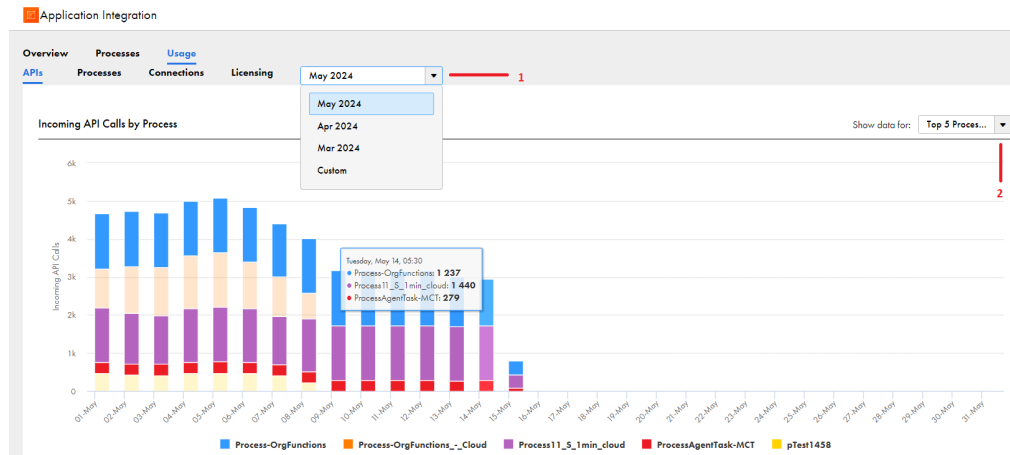
Hover over a bar or a point in the trend line of the chart to view the number of daily and cumulative incoming API calls.

Incoming API Calls by Process

The **Incoming API Calls by Process** panel shows the number of daily incoming API calls for a selected process or process category. By default, the panel shows information for the current month and the top 5 processes with the most incoming API calls.

You can use the **Incoming API Calls by Process** panel to find important processes that might need better optimization than other processes.

The following image shows the **Incoming API Calls by Process** panel:



1. From the list, select the month for which you want to view the daily incoming API calls. You can select the current month, one of the preceding 2 months, or a custom date range. If you select the **Custom** option, you can select a duration of 1 month in the past 13 months.

2. From the list, select a process or process category for which you want to view the daily incoming API calls. You can select one of the following values:

- All Processes
- Top 5 Processes
- Bottom 5 Processes
- A specific process

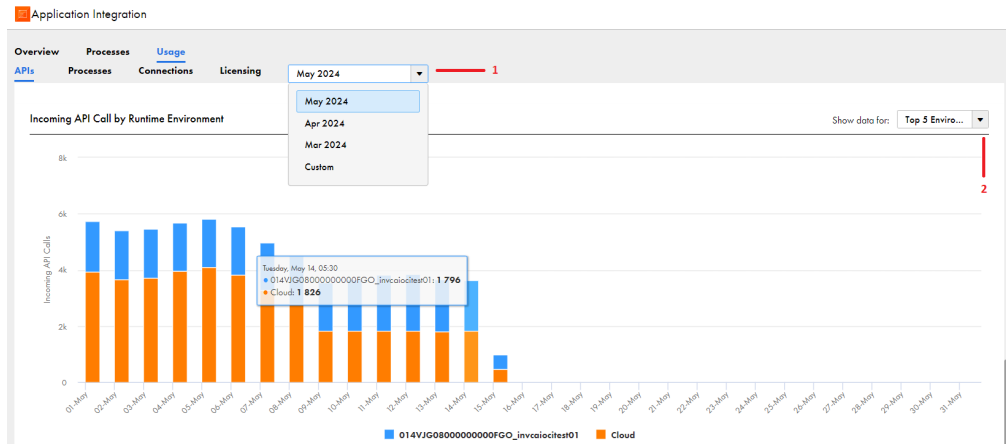
Hover over a bar in the chart to view the number of daily incoming API calls for the specified process or process category.

Incoming API Calls by Runtime Environment

The **Incoming API Calls by Runtime Environment** panel shows the number of daily incoming API calls for a selected runtime environment or runtime environment category. By default, the panel shows information for the current month and the top 5 runtime environments with the most incoming API calls.

You can use the **Incoming API Calls by Runtime Environment** panel to assess the load on different runtime environments and perform load balancing or increase the Secure Agent capacity.

The following image shows the **Incoming API Calls by Runtime Environment** panel:



1. From the list, select the month for which you want to view the daily incoming API calls. You can select the current month, one of the preceding 2 months, or a custom date range. If you select the **Custom** option, you can select a duration of 1 month in the past 13 months.

2. From the list, select a runtime environment or runtime environment category for which you want to view the daily incoming API calls. You can select one of the following values:

- All Environments
- Top 5 Environments
- Bottom 5 Environments
- A specific runtime environment
- Cloud Only
- All Agents

Hover over a bar in the chart to view the number of daily incoming API calls by runtime environment.

View Application Integration process runs

Click the **Processes** tab under the **Usage** tab to view analytics related to process runs. You can view daily and cumulative process runs, the number of runs for specific processes, and the number of processes that were run in a specific runtime environment.

The panels on the **Processes** tab show metrics for both parent processes and subprocesses. The panels exclude processes that were suspended at any point of time in their lifecycle.


You can use the following panels to view process run analytics:

Daily and Cumulative Process Runs

The **Daily and Cumulative Process Runs** panel shows the number of daily and cumulative process and subprocess runs. By default, the panel shows information for the current month and all runtime environments.

You can use the **Daily and Cumulative Process Runs** panel to assess whether there is a surge or dip in the number of process and subprocess runs, and accordingly plan for peak traffic.

The following image shows the **Daily and Cumulative Process Runs** panel:



1. From the list, select the month for which you want to view the for process runs. You can select the current month, one of the preceding 2 months, or a custom date range. If you select the **Custom** option, you can select a duration of 1 month in the past 13 months.

2. From the list, select the runtime environment or runtime environment category for which you want to view the process runs. You can select one of the following values:

- All Environments
- A specific runtime environment
- Cloud Only
- All Agents


Hover over a bar or a point in the trend line of the chart to view the number of daily and cumulative process runs.

Process Runs Distribution

The **Process Runs Distribution** panel shows the number of daily process and subprocess runs for a selected process or process category. By default, the panel shows information for the current month and the top 5 processes with the most runs.

You can use the **Process Runs Distribution** panel to find important processes that might need better optimization than other processes.

The following image shows the **Process Runs Distribution** panel:



1. From the list, select the month for which you want to view the number of daily process runs. You can select the current month, one of the preceding 2 months, or a custom date range. If you select the **Custom** option, you can select a duration of 1 month in the past 13 months.

2. From the list, select a process or process category for which you want to view the number of daily process runs. You can select one of the following values:

- All Processes
- Top 5 Processes
- Bottom 5 Processes
- A specific process


Hover over a bar in the chart to view the number of daily process runs for the specified process or process category.

Process Runs by Runtime Environment

The **Process Runs by Runtime Environment** panel shows the number of daily process and subprocess runs for a selected runtime environment or runtime environment category. By default, the panel shows information for the current month and the top 5 runtime environments with the most process runs.

You can use the **Process Runs by Runtime Environment** panel to assess the load on different runtime environments and perform load balancing or increase the Secure Agent capacity.

The following image shows the **Process Runs by Runtime Environment** panel:



1. From the list, select the month for which you want to view the number of daily process runs. You can select the current month, one of the preceding 2 months, or a custom date range. If you select the **Custom** option, you can select a duration of 1 month in the past 13 months.

2. From the list, select a runtime environment or runtime environment category for which you want to view the number of daily process runs. You can select one of the following values:

- All Environments
- Top 5 Environments
- Bottom 5 Environments
- A specific runtime environment
- Cloud Only
- All Agents

Hover over a bar in the chart to view the number of daily process runs by runtime environment.

View Application Integration connection calls

Click the **Connections** tab under the **Usage** tab to view analytics related to connections. You can view the number of daily and cumulative connections calls made, and also analyze the number of connection calls made for different endpoints, connection types, and runtime environments.


You can use the following panels to view connection call analytics:

Daily and Cumulative Connection Calls

The **Daily and Cumulative Connection Calls** panel shows the number of daily and cumulative connection calls. By default, the panel shows information for the current month and all runtime environments.

You can use the **Daily and Cumulative Connection Calls** panel to gain insights into connection usage.

The following image shows the **Daily and Cumulative Connection Calls** panel:



1. From the list, select the month for which you want to view the number of daily and cumulative connection calls. You can select the current month, one of the preceding 2 months, or a custom date range. If you select the **Custom** option, you can select a duration of 1 month in the past 13 months.

2. From the list, select the runtime environment or runtime environment category for which you want to view the number of daily and cumulative connection calls. You can select one of the following values:

- All Environments
- A specific runtime environment
- Cloud Only
- All Agents


Hover over a bar or a point in the trend line of the chart to view the number of daily and cumulative connection calls.

Connection Calls by Endpoints

The **Connection Calls by Endpoints** panel shows the number of daily connection calls for a selected endpoint or endpoint category. By default, the panel shows information for the current month and the top 5 endpoints in your organization.

You can use the **Connection Calls by Endpoints** panel to find the most commonly used connection endpoints.

The following image shows the **Connection Calls by Endpoints** panel:



1. From the list, select the month for which you want to view the number of daily connection calls. You can select the current month, one of the preceding 2 months, or a custom date range. If you select the **Custom** option, you can select a duration of 1 month in the past 13 months.

2. From the list, select an endpoint or endpoint category for which you want to view the number of daily connection calls. You can select one of the following values:

- All Endpoints
- Top 5 Endpoints
- Bottom 5 Endpoints
- A specific endpoint


Hover over a bar in the chart to view the number of daily process runs for the specified endpoint or endpoint category.

Connection Calls by Type of Connection

The **Connection Calls by Type of Connection** panel shows the number of daily connection calls for a selected connection type. By default, the panel shows information for the current month and the top 5 connection types with the most connection calls.

You can use the **Connection Calls by Type of Connection** panel to find the most commonly used connection types. It helps you assess whether there are more connection calls made to on-premises systems or Cloud-based systems.

The following image shows the **Connection Calls by Type of Connection** panel:



1. From the list, select the month for which you want to view the number of daily connection calls. You can select the current month, one of the preceding 2 months, or a custom date range. If you select the **Custom** option, you can select a duration of 1 month in the past 13 months.

2. From the list, select a connection type for which you want to view the number of daily connection calls. You can select one of the following values:

- All Types
- Top 5 Types
- Bottom 5 Types
- A specific connection type

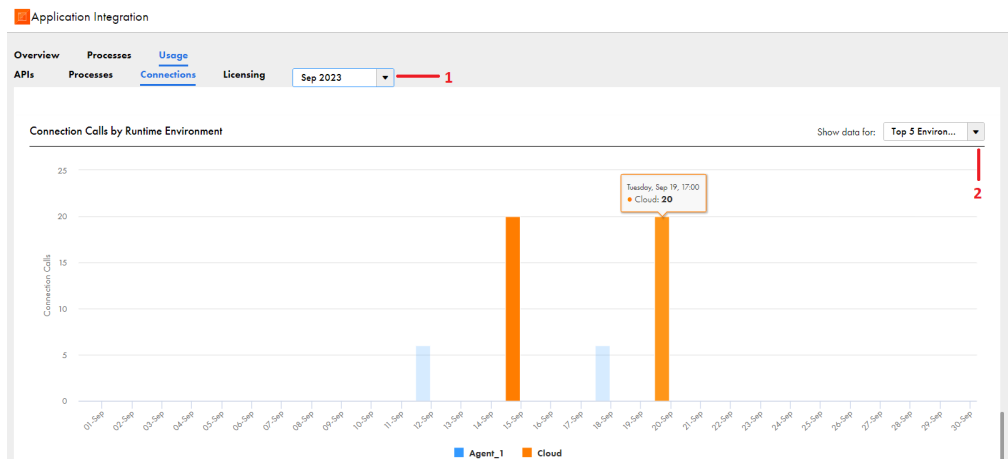
Hover over a bar in the chart to view the number of daily connection calls by connection type.

Connection Calls by Runtime Environment

The **Connection Calls by Runtime Environment** panel shows the number of daily connection calls for a selected runtime environment or runtime environment category. By default, the panel shows information for the current month and the top 5 runtime environments with the most connection calls.

You can use the **Connection Calls by Runtime Environment** panel to assess the load on different runtime environments and decide whether load balancing is needed.

The following image shows the **Connection Calls by Runtime Environment** panel:



1. From the list, select the month for which you want to view the number of daily connection calls. You can select the current month, one of the preceding 2 months, or a custom date range. If you select the **Custom** option, you can select a duration of 1 month in the past 13 months.

2. From the list, select a runtime environment or runtime environment category for which you want to view the number of daily connection calls. You can select one of the following values:

- All Environments
- Top 5 Environments
- Bottom 5 Environments
- A specific runtime environment
- Cloud Only
- All Agents


Hover over a bar in the chart to view the number of daily connection calls by runtime environment.

View API transactions against licensing limit

Click the **Licensing** tab under the **Usage** tab to assess whether the total number of API transactions for all runtime environments in your organization falls within or exceeds the maximum API transactions count defined in Administrator for your Application Integration license.

If the total number of API transactions exceeds the defined maximum API transactions count, the chart indicates the difference as a percentage value. You might have cost implications if you exceed the allowed maximum limit.

In the following image, the chart shows the total number of API transactions for each runtime environment and indicates that the total number of API transactions as on March 19, 2020 has exceeded the licensing limit by 343.4 percent:




1. From the list, select the month for which you want to view the number of API transactions. You can select the current month, one of the preceding 2 months, or a custom date range. If you select the **Custom** option, you can select a duration of 1 month in the past 13 months.

You can use the following panel to view the daily and cumulative incoming API calls against your API transactions limit:

Daily and Cumulative Incoming API Calls with Trendline

The **Daily and Cumulative Incoming API Calls with Trendline** panel shows the number of daily and cumulative incoming API calls for a selected runtime environment or runtime environment category. The count also includes the incoming API calls from scheduled processes and event-based processes. The **Daily and Cumulative Incoming API Calls with Trendline** panel helps you compare the total number of API calls against the maximum API transactions count defined in Administrator for your Application Integration license. By default, the panel shows information for the current month and all runtime environments.

The following image shows the **Daily and Cumulative Incoming API Calls with Trendline** panel:



1. From the list, select the runtime environment or runtime environment category for which you want to view the number of daily and cumulative incoming API calls. You can select one of the following values:

- All Environments
- A specific runtime environment
- Cloud Only
- All Agents

To view the number of daily and cumulative incoming API calls for a specific month, scroll up and select the month from the list. You can select the current month, one of the preceding 2 months, or a custom date range. If you select the **Custom** option, you can select a duration of 1 month in the past 13 months.

Hover over a bar or a point in the trend line of the chart to view the number of daily and cumulative incoming API calls and compare it with your API transactions limit. The red line in the chart indicates the maximum API transactions count defined in Administrator for your Application Integration license. Points and bars that extend above the red line are outliers.

CHAPTER 6

Monitor Informatica Intelligent Cloud Services Data Profiling

If your organization uses Data Profiling, you can use Operational Insights to view the job status for data profiling tasks.

For more information about using Data Profiling, see the Data Profiling help.

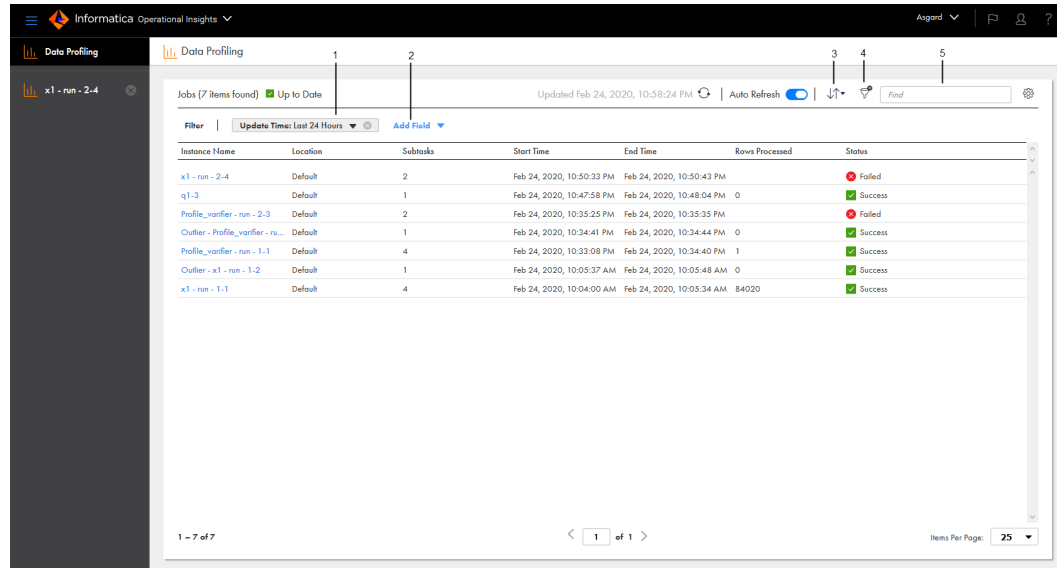
To monitor data profiling jobs in Operational Insights, you need access to Data Profiling and Operational Insights. You also need Operational Insights - view privilege.

For more information or to request this feature, contact Informatica Global Customer Support.

View Data Profiling service jobs

To view the details about the jobs that have run in your organization, click **Data Profiling** in the left hand side navigation bar.

The following image shows the **Data Profiling** page:



1. Change the time period. You can view jobs run in the last 24 hours, in the last week, in the last month, in the last year, or you can enter a custom range.
2. Add a new filter.
3. Sort the jobs on the page.
4. Add or remove filters.
5. Search for jobs.

By default, the **Jobs** page displays jobs that completed in the last 24 hours. You can apply filters to view jobs that were completed in the last 30 days. You can apply the following filters:

- Instance Name
- Asset Name
- Location
- Subtasks
- Runtime Environment
- Start Time
- End Time
- Duration
- Rows Processed
- Started By
- Status

By default, the following properties are displayed for each job:

- Instance Name
- Location
- Subtasks

- Start Time
- End Time
- Rows Processed
- Status

You can also display the following additional properties by right-clicking in the column heading area:

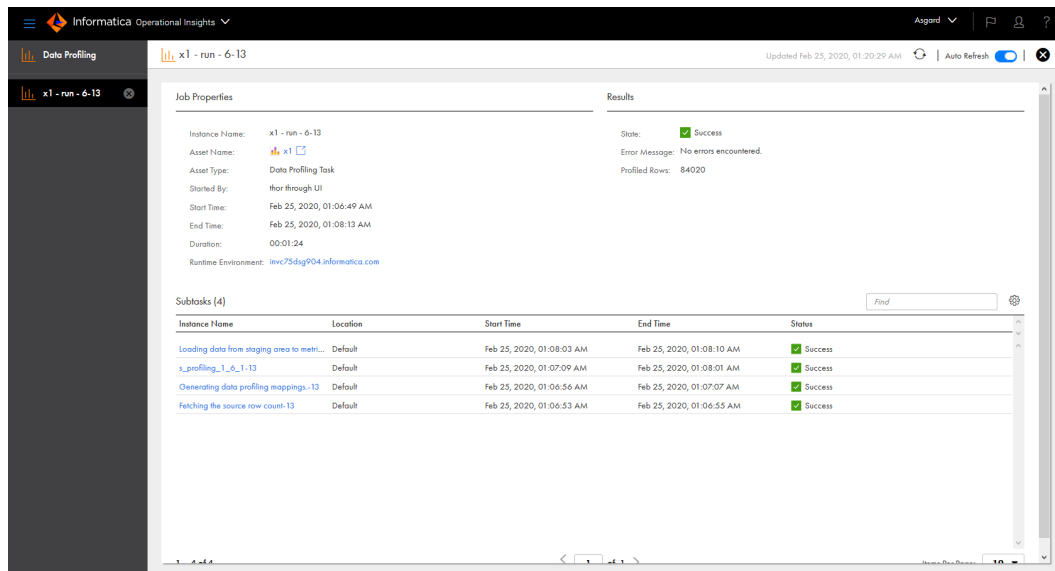
- Asset Name
- Runtime Environment
- Secure Agent
- Update Time
- Duration
- Started By
- Error Message

View details for a specific job in data profiling

You can drill down on a specific job to view the job details and subtasks of the job.

Click the instance name in the **Instance Name** column to open the **Job details** page.

The following image shows the job details page for a data profiling task:



You can view the subtasks for the instance in the **Subtasks** area. Click a subtask to view the subtask details.

You can view the runtime environment and the Secure Agent for the following subtasks in Data Profiling, Monitor, and Operational Insights:

- Fetching the source row count
- s_profiling

- Drilldown
- Query

Note: The Runtime Environment field displays the name of the Secure Agent Group.

Click the **Download Session Log** to download the session log file. You can view the following details for the Secure Agent in the session log file for the profile mapping jobs:

- Task Name. The name of the profiling task.
- Agent Group Id. The ID of the Secure Agent Group.
- Agent Group Name. The name of the Secure Agent Group.
- Agent Id. The ID of the Secure Agent.
- Agent Name. The name of the Secure Agent.

CHAPTER 7

Monitor Informatica Intelligent Cloud Services Data Ingestion and Replication

You can monitor the progress, performance, and status of Data Ingestion and Replication jobs from the following services:

Monitor

Monitor all jobs or all running jobs, including Data Ingestion and Replication jobs and Data Integration jobs, on the **All Jobs** and **Running Jobs** pages in Monitor. You can use this interface to do the following tasks:

- View general job properties such as the instance name, project location, subtasks, start time, end time, and status.
- Perform some actions on a job, depending on the job status and task type, using the **Actions** menu at the right end of each job row.
- Drill down on a job to see job details.

Operational Insights

Monitor ingestion jobs that any member of your organization deployed on the **Data Ingestion and Replication** page in Operational Insights. The **Data Ingestion and Replication** page has the following tabs:

- **Overview** tab. Includes buttons that you can use to filter the list of ingestion jobs by job type and status. Each button displays a summary count of the jobs that match that filter. There's also an option for showing the jobs that have statuses requiring attention, for example, failed jobs or jobs running with errors or warnings.
- **All Jobs** tab. Lists all types of ingestion and replication jobs that any member in your organization created and deployed. To help you find a job, the tab includes **Filter**, **Find**, and **Sort** features. Also, each job row contains an **Actions** menu to perform actions on the job, such as Stop or Redeploy, depending on the job type and status.

Both tabs show general job properties, including the instance name, project location, task type, runtime environment, start time, duration, and status. From either tab, you can click a job name to drill down to see job details.

Data Integration

Monitor the ingestion and replication jobs and data integration jobs for the tasks that you deployed on the Data Integration unified Home page and on the **My Jobs** page accessed from the Home page's left navigation bar. You can view the following types of information:

- The **Recent Jobs** panel on the Home page shows the status of your five most recent jobs. Use this panel to quickly check the job status. Click the job name to drill down to see job details. Click **View All** to go to the **My Jobs** page to see all of your jobs.
- The **My Jobs** page lists all of your jobs. You can use the **Filter**, **Find**, and **Sort** features to find the job you need. Use the **Actions** menu at the right end of each job row to perform some actions on the job, depending on the job status and task type. You can click a job name to drill down on a job to see job details.

Note: Usually, a job name corresponds to the ingestion and replication task name. On the **All Jobs** and **Running Jobs** pages in Monitor, on the **My Jobs** page Data Integration, and in the **Recent Jobs** panel on the unified Home page, the job name has the following format:

```
<taskname>-<job_instance_number>
```

The instance number is incremented each time the job is deployed.

Monitoring your ingestion and replication jobs

You can monitor the Data Ingestion and Replication jobs for the tasks that you deployed in the Data Integration, Monitor, and Operational Insights services.

You can monitor Data Ingestion and Replication jobs from the unified Home page in Data Integration in the following locations:

- In the navigation bar on the left-hand side, click **My Jobs** to view all your jobs, including Data Ingestion and Replication and Data Integration jobs. This page shows information for each job, such as the instance name, location, start time, end time, rows processed, and status. To drill down details of a specific job, click the instance name and a page for the job appears. The details vary by the type of the ingestion job. To find a job on the **My Jobs** page, use any one of the following methods:
 - To *sort* the listed jobs, click a column heading or click the **Sort** icon and select a field to sort by.
 - To *find* a job based on the job name, enter the job instance name, or any part of the name, in the **Find** text box.
 - To *filter* the list of jobs, click the **Filter** icon. Then, click **Add Filter** and enter filter criteria for one or more of the listed fields.

To change the columns displayed on **My Jobs** page, right-click on the column header row and change the field selections.

From the **Actions** menu at the right end of each job row, you can perform some actions on the job, depending on the job status and task type.

In the **My Jobs** page, the retention time for undeployed jobs is based on the Job Log Service (JLS) purge policy.

- In the **Recent jobs** panel on the unified Home page, view the jobs that you recently deployed. You can use this panel to quickly check the job statuses.

Note: When you deploy an application ingestion and replication task or database ingestion and replication task, the newest job doesn't automatically appear at the top of the jobs list on the **My Jobs** page. You can sort jobs by using the **Start Time** property or use the find and filter features to search for a specific job.

You can also navigate to your jobs in the lists of **All Jobs** and **Running Jobs** pages in Monitor and in the **All Jobs** tab on the Data Ingestion and Replication page in Operational Insights.

For descriptions of the columns in Monitor and Operational Insights, see [“Job properties” on page 77](#).

Monitoring all ingestion and replication jobs

In Operational Insights and Monitor, you can monitor all Data Ingestion and Replication jobs that were deployed, including application ingestion and replication jobs, database ingestion and replication jobs, file ingestion and replication jobs, and streaming ingestion and replication jobs.

Monitoring jobs in Operational Insights

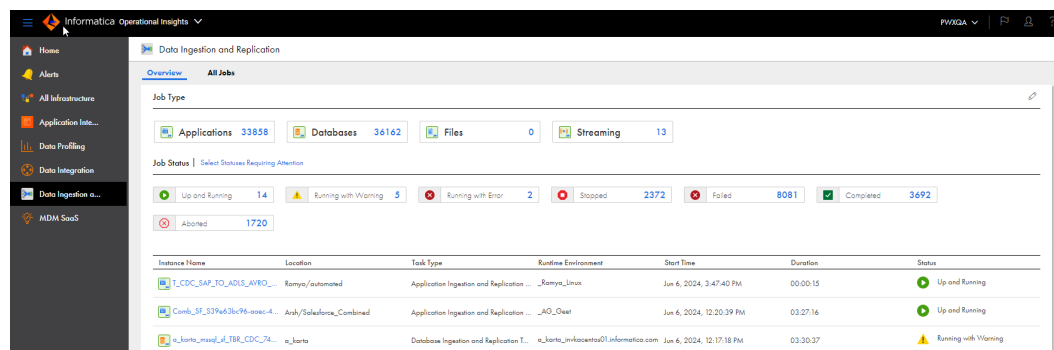
The Data Ingestion and Replication page in Operational Insights displays only Data Ingestion and Replication jobs.

The Data Ingestion and Replication page has the following tabs:

- The **Overview** tab displays buttons that you can use to filter the list of ingestion and replication jobs by job type and status.
- The **All Jobs** tab lists all types of ingestion and replication jobs that any member in your organization created and deployed.

Overview tab

The **Overview** tab initially lists all types of ingestion and replication jobs with any status by default. Each button shows the number of jobs that matches the job type or status. For example:




You can use the buttons at the top to filter the jobs by job type and status. If you want to see jobs of a particular job type and status, first select the job type and then select the status. You cannot select multiple status buttons at the same time. To clear a filter, click the selected button again.

To control the status buttons that appear on the **Overview** tab, click the **Edit** (pencil) icon. Then in the **Reorder Job Status** dialog box, select the **Visibility** check box next to each job status for which you want to display buttons and jobs.

To rearrange the order of the job status buttons, click the **Edit** (pencil) icon. Then in the **Reorder Job Status** dialog box, select and drag a job status row up or down.

To show only the jobs that have a status of concern, click **Select Statuses Requiring Attention**. The following example image shows all jobs with statuses requiring attention:

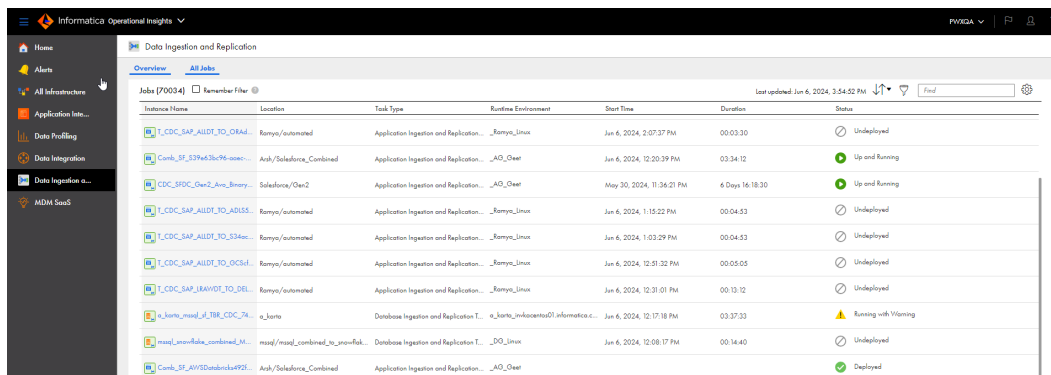


This option lists the application ingestion and replication jobs or database ingestion and replication jobs with the **Failed** or **Running with Warning** status, file ingestion and replication jobs with the **Failed** status, and streaming ingestion and replication jobs with the **Running with Error** or **Running with Warning** status. To clear the filter, click **Select Statuses Requiring Attention** again. If you want to see jobs that require attention for a specific job type, first select the job type and then click **Select Statuses Requiring Attention**.

Note: All filters that you set on the **Overview** tab or in the **Reorder Job Status** dialog box are active only for the current session or until you change them during the session.

All Jobs tab

The **All Jobs** tab lists all Data Ingestion and Replication jobs for your organization. For example:



If the list of jobs is long, you can use any of the following methods to find a job:

- To **find** a job based on the job name, enter the job instance name, or any part of the name, in the **Find** text box. With a partial name, the Find operation looks for that particular string anywhere in the instance name. You can include the percent sign (%) wildcard within an instance name search string to represent one or more characters, such as "ing2%798". Do not include the following symbols: question mark (?), number sign (#), or ampersand (&). If you include any of these symbols, the Find operation returns no results.
- To **filter** the list of jobs, click the **Filter** icon. Then click **Add Filter** and enter filter criteria for one or more fields. For the **Instance Name** field, you can enter the full job instance name or part of the name. You can include the percent sign (%) wildcard in the instance name value to represent one or more characters within the name, for example, "vp%test3". Your filter is saved for your user name only, for the current session until you change it. You can save the filter for the subsequent sessions by selecting the **Remember Filter** check box. To clear existing filter criteria, click the **Filter** icon again.
- To **sort** the listed jobs, click a column heading or click the **Sort** up or down arrows and select a field to sort by. The default sort order for application ingestion and replication jobs, database ingestion and replication jobs, and streaming ingestion and replication jobs is the time of task deployment, from latest to earliest. The default sort order for file ingestion and replication jobs is the job start time, from latest to earliest.

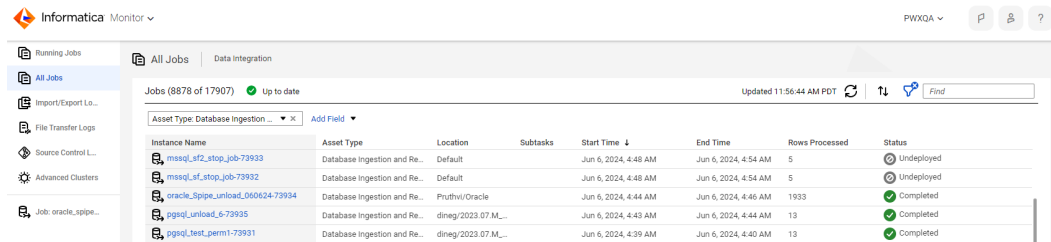
You can also use the Actions menu at the end of each row to perform an action on the job, such as Redeploy, depending on the job type and status.

Tip: To change the line spacing in the list, click the **Settings** icon to the right of the **Find** box.

Monitoring jobs in Monitor

The **All Jobs** and **Running Jobs** pages in Monitor shows all jobs and all running jobs respectively, including Data Ingestion and Replication jobs and Data Integration jobs if you run both types of jobs.

The following image shows a sample **All Jobs** page in Monitor:



| Instance Name | Asset Type | Location | Subtasks | Start Time | End Time | Rows Processed | Status |
|-----------------------------------|------------------------------|--------------------|----------|----------------------|----------------------|----------------|------------|
| messq_sf2_stop_job-73933 | Database Ingestion and Re... | Default | | Jun 6, 2024, 4:48 AM | Jun 6, 2024, 4:54 AM | 5 | Undeployed |
| messq_sf2_stop_job-73932 | Database Ingestion and Re... | Default | | Jun 6, 2024, 4:48 AM | Jun 6, 2024, 4:54 AM | 5 | Undeployed |
| oracle_spispe_unload_060624-73934 | Database Ingestion and Re... | Pruthvi/Oracle | | Jun 6, 2024, 4:44 AM | Jun 6, 2024, 4:46 AM | 1933 | Completed |
| pssqj_unload_6-73935 | Database Ingestion and Re... | dimeg/2023.07.M... | | Jun 6, 2024, 4:43 AM | Jun 6, 2024, 4:44 AM | 13 | Completed |
| pssqj_test_perm1-73931 | Database Ingestion and Re... | dimeg/2023.07.M... | | Jun 6, 2024, 4:39 AM | Jun 6, 2024, 4:40 AM | 13 | Completed |

You can monitor all running jobs on the **Running Jobs** page. The **Running Jobs** page lists the jobs that are starting, queued, running, and suspended. To drill down details of a specific job from either **All Jobs** or **Running Jobs**, click the instance name and a page for the job appears.

When you deploy an application ingestion and replication or database ingestion and replication, the newest job doesn't automatically appear at the top of the jobs list on the **All Jobs** page. You can sort jobs by using the **Start Time** property or use the *find* and *filter* features to search for a specific job.

To find a job on the **All Jobs** or **Running Jobs** pages, use any one of the following methods:

- To *sort* the listed jobs, click a column heading or click the **Sort** icon and select a field to sort by.
- To *find* a job based on the job name, enter the job instance name, or any part of the name, in the **Find** text box.
- To change the columns displayed on either pages, right-click on the column header row and change the field selections.
- To *filter* the list of jobs on the **All Jobs** page, click the **Filter** icon. Then, click **Add Filter** and enter filter criteria for one or more of the listed fields.

Use the **Actions** menu at the end of each job row to perform actions on a job based on the job status and task type. You can also stop a Data Ingestion and Replication job on the **All Jobs** and **Running Jobs** pages using the **Stop** icon beside the **Actions** menu.

In the **All Jobs** page, the retention time for undeployed jobs is based on the Job Log Service (JLS) purge policy.

Job properties

The lists of ingestion and replication jobs on the **My Jobs** page and on Monitor and Operational Insights pages display properties for each job, which you can review to get a high-level view of the job statuses.

The following table describes the job properties shown on the **Overview** and **All Jobs** tabs of the **Data Ingestion and Replication** page in Operational Insights:

| Property | Description |
|---------------------|---|
| Instance Name | <p>The generated name of the job instance in the following format: <code><taskname>-<job_instance_number></code></p> <p>You can click the instance name to view detailed information about the job. Note: If you edit the name of the associated ingestion and replication task, the job name remains the same.</p> |
| Location | <p>The project or project\subfolder, where the task definition associated with the job exists. For example: <code>Myproject\Oracle</code></p> <p>If you move a task definition to another folder, the Location value is not updated.</p> |
| Task Type | <p>The type of ingestion and replication task. The type can be one of the following values:</p> <ul style="list-style-type: none"> - Application Ingestion and Replication Task - Database Ingestion and Replication Task - File Ingestion and Replication Task - Streaming Ingestion and Replication Task |
| Runtime Environment | <p>The name of the runtime environment that includes the Secure Agent on which the job runs.</p> |
| Start Time | <p>For application ingestion and replication and database ingestion and replication jobs, the date and time when the job was deployed.</p> <p>For file ingestion and replication jobs, the date and time when the job started.</p> <p>For streaming ingestion and replication jobs, the date and time when the job was deployed.</p> |
| Duration | <p>For application ingestion and replication and database ingestion and replication jobs, the amount of time that the job has run since it was deployed. For jobs that have the Completed, Stopped, Failed, or Aborted status, the amount of time between the date and time the job was deployed and when it acquired its current status.</p> <p>For file ingestion and replication jobs, the amount of time that the job has run.</p> <p>For streaming ingestion and replication jobs, the amount of time that the job has been running.</p> |
| Status | <p>The current status of the job, such as Deploying, Up and Running, or Undeployed.</p> <p>The set of valid statuses vary by type of ingestion and replication task. For more information, see the "Job Overview tab" section in "Application ingestion and replication job details" on page 79, "Database ingestion and replication job details" on page 85, "File ingestion and replication job details" on page 92 ("Results" section), or "Streaming ingestion and replication job details" on page 95.</p> |

The following table describes the job properties shown on the **All Jobs** and **Running Jobs** pages in Monitor and on the **My Jobs** page accessed from the Data Integration Home page:

| Property | Description |
|----------------|---|
| Instance Name | <p>The generated name of the job instance in the following format: <taskname>-<job_instance_number></p> <p>You can click the instance name to view detailed information about the job.</p> <p>Note:</p> <ul style="list-style-type: none"> - If you edit the name of the associated ingestion and replication task, the job name remains the same. - The job name contains an underscore (_) instead of a hyphen (-): <taskname>_<job_instance_number> |
| Location | <p>The project or project\subfolder, where the task definition associated with the job exists. For example: Myproject\Oracle</p> <p>If you move a task definition to another folder, the Location value is not updated.</p> |
| Subtasks | <p>The number of subtasks on source objects.</p> <p>Note: For application ingestion and replication jobs and database ingestion and replication jobs, the subtask count is not displayed.</p> |
| Start Time | <p>For application ingestion and replication and database ingestion and replication jobs, the date and time when the job was deployed.</p> <p>For file ingestion and replication jobs, the date and time when the job started.</p> <p>For streaming ingestion and replication jobs, the date and time when the job was deployed.</p> |
| End Time | <p>For application ingestion and replication and database ingestion and replication jobs, the date and time when the job ended because it completed processing, was stopped, or failed. This field is not displayed for running jobs.</p> |
| Rows Processed | <p>For application ingestion and replication jobs and database ingestion and replication jobs, the total number of rows that the job has processed and successfully written to the target. If a job encounters an error during row processing, the job status switches to Failed and the rows with the error are not included in the Rows Processed count.</p> <p>For file ingestion and replication jobs, the total number of files that the job has processed and successfully written to the target. If a job encounters an error during file processing, the job status switches to Failed and the files with the error are added in the Failure Rows count. In this case, the Rows Processed property displays the total number of files processed irrespective of their statuses.</p> <p>For streaming ingestion and replication jobs, the total number of events that the job has processed and successfully written to the target. If a job encounters an error during event processing, the job status switches to Running with Warning or Running with Error and the events associated with the error are not included in the Rows Processed count.</p> |
| Status | <p>The current status of the job, such as Deploying, Up and Running, or Undeployed.</p> <p>The set of valid statuses vary by type of ingestion and replication task. For more information, see the "Job Overview tab" section in "Application ingestion and replication job details" on page 79, "Database ingestion and replication job details" on page 85, "File ingestion and replication job details" on page 92 ("Results" section), or "Streaming ingestion and replication job details" on page 95.</p> |

Note: You can change the columns displayed on the **My Jobs** page or on the **All Jobs** and **Running Jobs** pages in Monitor by right-clicking in the column header row and changing the field selections. If you add the **Success Rows** and **Failure Rows** fields but your ingestion and replication task type does not support these fields, a -1 is displayed, which you can ignore.

Viewing details for an ingestion and replication job

You can drill down on a specific Data Ingestion and Replication job from any of the following locations:

- The **My Jobs** page that's accessed from the navigation bar on the Data Integration Home page
- The **All Jobs** and **Running Jobs** pages in Monitor
- The **Overview** and **All Jobs** tabs on the **Data Ingestion and Replication** page in Operational Insights

To view job details, click the job name in the jobs list. A page for the job appears. The details vary by type of ingestion job.

Application ingestion and replication job details

You can drill down on a job to see job details from the **My Jobs** page accessed from the Data Integration Home page, from the **All Jobs** and **Running Jobs** pages in Monitor, and from the **Overview** and **All Jobs** tabs on the **Data Ingestion and Replication** page in Operational Insights.

After you click a job name, job information is displayed in the **Task Summary**, **Overview**, **Object Detail**, and **Alerts** panes. The **Alerts** pane is displayed only for incremental load and combined initial and incremental load jobs in Operational Insights.

Tip: Click the expander arrow next to the pane name to expand or collapse the pane.

Task Summary

View a diagram that shows the source and target connection types and connection names. It also shows the calculated data throughput, in rows per second if the job has successfully replicated data to the target, regardless of the job's current status. If the calculated value is 0, indicating no data has flowed to the target, the throughput is not displayed.

In the upper right corner, the **Status** field displays the job status, which can be one of the following values:

- **Up and Running.** The job is running.
- **Running with Warning.** The job is running with a warning. This state can also occur when one or more table-specific subtasks fail but some subtasks are still running.
- **On Hold.** The job is in a paused state while the Database Ingestion and Replication (DBMI) agent is being updated.
- **Stopping.** The job is stopping in response to a Stop request.
- **Stopped.** The job was intentionally stopped.
- **Failed.** The job ended abnormally, the task deployment to the job failed, or one or more table-specific subtasks failed. Also, for an initial load job, the job was stopped.
- **Deploying.** The job is being deployed.
- **Deployed.** The job has been deployed.
- **Aborting.** The job is stopping immediately in response to an Abort request.
- **Aborted.** The job has been aborted.
- **Undeploying.** The job is being undeployed.
- **Undeployed.** The job has been undeployed.
- **Completed.** The job completed successfully.

Also, for incremental load jobs and combined initial and incremental load jobs, you can download the job execution log for the entire job run. In the *Select log* list under the diagram, select one of the following log types:

- **Complete Log.** The entire log, including all types of messages. It is available for any job that ran, regardless of its state.
- **Error.** The error log, which includes messages only for errors that occurred. It is available for Failed jobs only. Use this log to determine the reason for the job failure, for example, the deployment failed. If the log file ends with an ellipsis (...), the log has been truncated because of its long length. In this case, download the Complete Log to see all error messages.


Note: For initial load jobs, you can get the job log for a specific source object from the **Object Detail** tab.

Overview

View information about the job, including the runtime environment, task name, load type, task type, task definition location, user who started the job, records read, records written, number of subtasks, and schedule name, if a schedule was used.

The circle image displays the number of subtasks on source tables by status. The color of the circle's rim corresponds to the status.

The following image shows a sample **Overview** pane in Monitor:



The following table describes the summary job properties and statistics:

| Property | Description |
|----------------------|--|
| Duration (on left) | The amount of time, in the hh:mm:ss format, that the job ran before it ended. |
| Start Time (on left) | The date and time when the job was deployed. |
| End Time (on left) | The date and time when the job ended because it completed processing, was stopped, or failed. This field is not displayed for running jobs |
| Runtime Environment | The name of the runtime environment that the job uses to run. |

| Property | Description |
|-----------------|---|
| Agent Name | <p>For incremental load and combined initial and incremental load jobs, the name of the Secure Agent that the job uses to run.</p> <p>In case of a failover from one Secure Agent to another, Application Ingestion and Replication displays the latest Secure Agent assigned to the job.</p> <p>Note: For initial load jobs, the Agent Name is displayed in the Object Detail pane. For combined initial and incremental load jobs, the Agent Name is displayed both in the and the Object Detail panes.</p> |
| Task Name | <p>The name of the associated application ingestion and replication task. You can click the task-name link to view or edit task details in Data Ingestion and Replication, if necessary. If you edit the task, you must redeploy it for the updated task definition to be used for a job.</p> |
| Load Type | <p>The type of load operation that the job performs. Options are:</p> <ul style="list-style-type: none"> - Initial Load. Loads a snapshot of source data read at a specific point-in-time to a target. - Incremental Load. Loads incremental data changes to a target on a continuous basis, until the job is stopped or ends. - Initial and Incremental Load. Performs an initial load and then automatically switches to an incremental load. |
| Task Type | <p>The type of task, which is Application Ingestion and ReplicationTask.</p> |
| Task Location | <p>The project or project folder that contains the application ingestion and replication task definition.</p> |
| Started By | <p>The name of the user who started the job.</p> |
| Records Read | <p>The number of records that were read from the source.</p> <p>Note: For a combined initial and incremental load job, the Records Read count might be greater than the total number of object-level DML change records written. This behavior occurs because the initial load or resync processing always starts after change data capture has begun. As a result, some change records are included in the Records Read count and then discarded before initial load or resync processing starts. These discarded records cause the Records Write count to be less than the Records Read count.</p> |
| Records Written | <p>The number of records that were successfully replicated to the target.</p> <p>Note: The Records Written value might be different from the Records Read value if source records are discarded. For example, in a combined initial and incremental load job, change records read from the source before the initial unload phase completes are discarded because they're not yet needed.</p> |
| Subtasks | <p>The number of subtasks that the application ingestion and replication job used to replicate data from source tables to the target. When a job runs, it uses a separate subtask to process each source table.</p> |
| Schedule | <p>For initial load jobs, the name of the schedule that is used to run the job or "No schedule" if you run the job manually.</p> |

Object Detail

On the **Object Detail** pane lists subtasks on the source tables. You can view statistics and status information by source table from the last run of a application ingestion and replication job. When you click the expander arrow next to an object name, counts of processed inserts, updates, deletes, and LOB changes are shown for the table.

The following image shows a sample **Object Detail** pane:

Object Detail (14)

| Object | Target Object | Agent Name | Records Read | Records Written | Task Duration | Status | Log |
|------------------------|----------------------------|------------|--------------|-----------------|---------------|-----------|------------|
| MSD365.annotations | PADMINI.AIN-1271_mrel_a... | N/A | 12 | 12 | 00:00:29 | Completed | Select log |
| MSD365.appmodules | PADMINI.AIN-1271_mrel_a... | N/A | 12 | 12 | 00:00:23 | Completed | Select log |
| MSD365.documenttemplat | PADMINI.AIN-1271_mrel_d | N/A | 0 | 0 | 00:00:27 | Completed | Select log |

Note: This pane shows information for the latest job run. This tab is blank for jobs that have not run or are resuming.

The following table describes the **Object Detail** fields that are displayed for each table, depending on the load type and status:

| Column | Description |
|-----------------|---|
| Object | <p>The name of the source table or view for which data was propagated to the target.</p> <p>For an incremental load job or a combined initial and incremental load job, click the arrow icon to the left of the object name to display detailed counts of Inserts, Deletes, Updates, LOBs, and DDL statements processed. For a combined initial and incremental load job, the Unload Count field is also displayed to show the number of records that the initial load portion of processing read from the source. The following usage notes apply to the detailed CDC counts:</p> <ul style="list-style-type: none"> - The counts are only for the current job run. If you stop and restart the job, the counts start over from zero. Do not use these counts to identify the number of rows written to the target. - The counts are based on rows read from the source and do not reflect the records written to the target. Target write operations might be optimized by combining operations and reducing the number of physical writes. In this case, the counts might not match the number of write operations. - The value N/A means that the count value is not applicable for the count type or the value has not yet been calculated. - The Unload Count might not reflect the number of source records at the time the job is started or resynchronized because of a delay in the start of unload processing. Between the time of the unload request and start of unload processing, rows might be added to or deleted from the source table. |
| Target Object | The name of the target object that is mapped to the source object. |
| Agent Name | <p>For initial load and combined initial and incremental load jobs, the name of the Secure Agent on which the job was run.</p> <p>For incremental load jobs, the Agent Name is displayed in the Overview section of the Task Summary pane. For combined initial and incremental load jobs, the Agent Name is displayed both in the Task Summary and the Object Detail panes.</p> |
| Records Read | For an initial load job, the number of records that were read from the source. For other load types, this information is available only at the job-level on the Job Overview tab. |
| Records Written | <p>For an initial load job, the number of records that were successfully written to the target. For other load types, this information is available only at the job-level on the Job Overview tab.</p> <p>Note: The Records Read value can be greater than the Records Written value if some records read from the source were discarded. For example, in a combined initial and incremental job, any source change records read before the initial unload phase of the job has completed are discarded.</p> |
| Task Duration | <p>For an initial load job, the amount of time the subtask that processed the source table ran before it completed or was stopped. For other load types, this information is available only at the job-level on the Job Overview tab.</p> <p>When a job runs, it uses a separate subtask to process each source table.</p> |

| Column | Description |
|--------|---|
| Stage | <p>For a combined initial and incremental load job, this column shows the stage in the transition from initial load processing to CDC processing for the table-specific job subtask. This column does not appear for other load types.</p> <p>The stage can be one of the following values:</p> <ul style="list-style-type: none"> - Not Started. Initial load processing has not yet started for the table, or if an error occurred and the table is in the Error on Retry state, the next attempt to process the table has not yet started. - Started. Initial load processing has started. - Unloading. The subtask is unloading data from the table as part of initial load processing. - Unloaded. The subtask has finished unloading data from the table as part of initial load processing. - Completed. The subtask completed initial load processing of the table. - Normal. The subtask completed initial load processing of the table and has started CDC processing of the table. - Cancelled. Initial load processing was cancelled or stopped. - Error. The subtask detected an error in the source table. <p>Actions menu > Resync</p> <p>For a subtask in a combined initial and incremental load job, if the subtask stage is Normal and the subtask status is any status other than Queued or Starting, the Actions menu is displayed on the right end of the subtask row. From the Actions menu, you can select Resync to resynchronize the source and target objects. For more information, see "Resynchronizing source and target objects" in Data Ingestion and Replication help.</p> |

| Column | Description |
|--------|---|
| Status | <p>The status of the job subtask for the source object. Note: If the job stops running, the subtask status reflects the status last collected before the job ended. For example, the job might be aborted but the subtask is in a Running status.</p> <p>The status can be one of the following values:</p> <ul style="list-style-type: none"> - Queued. The subtask has not yet started running. - Starting. The subtask is starting. - Started. For a combined initial and incremental load job, the subtask has started. - Running. The subtask is running. - On Hold. The subtask, as well as the job, is in a paused state while the Database Ingestion and Replication (DBMI) agent is being updated. - Completed. The subtask completed processing successfully. - Stopping. The subtask is stopping in response to a Stop request. - Stopped. The subtask has stopped. - Aborting. The subtask is ending immediately in response to an Abort request. - Aborted. The subtask has been aborted. - Failed. The subtask ended unexpectedly. - Error. The subtask is in error and no longer writing data to the target table. For a combined initial and incremental load job, the subtask might be running and processing incremental change data but no data is being sent to the target. - Error on Retry. An error occurred on the last retry of subtask processing, and now the subtask is waiting to retry processing again. <p>Note: If a DDL change occurs on a source table and then you resume the job, the table subtask state might not change as expected until the first DML operation occurs on the source table.</p> |
| Log | <p>You can download a job execution log for a source object. The type and availability of the log depends on the load type and status. Options are:</p> <ul style="list-style-type: none"> - Complete. The complete log for an object subtask from job execution. This log type is available for a Completed, Failed, or Aborted subtask in an initial load job. - Error. The log that contains error messages. This log type is available only for a Failed or Error subtask in an initial load or incremental load job. - Stage Log. The log that covers the transition from the initial processing phase to the incremental processing phase in a combined initial and incremental load job for a source object. <p>To download a log, click the Download icon.</p> <p>Note: If you undeployed the job, you can download the log for a table only if the associated task has not been deleted.</p> <p>For incremental load jobs, you can get the complete log and error log for the entire job run from the Task Summary pane.</p> |

Note: To control the line spacing in the list, click the **Settings** icon next to the *Find* box.

Alerts

The **Alerts** pane appears on the **Data Ingestion and Replication** page in Operational Insights for a selected incremental load or combined initial and incremental load job. On the **Alerts** pane, you can view alert messages that appear for certain events, such as source schema changes, during incremental load or combined initial and incremental load processing.

Note: The **Alerts** pane displays alert messages for all detected schema changes even if you set the schema drift options for the associated task to Ignore.

You can filter the list of alerts based on severity or a date range. To specify a date range, enter one of the following types of values in the **Filter** field:

- **Any Time** for all stored alerts.
- **Today** for alerts issued today from midnight to 11:59 pm.
- **Last Week, Last Month, or Last Year** to show alerts from the beginning of last week, month, or year to present.

- **Custom** to specify a custom date range that consists of a beginning date and time and an ending date and time.

The list of alerts includes the following columns:

| Column | Description |
|--------|---|
| Level | Severity level of the alert message, such as Critical or Warning. |
| Code | Alphanumeric code that identifies the alert type followed by the date and time when the event occurred. |

Click the expander arrow to display a description of the event.

You can configure alert notifications for application ingestion and replication jobs from the **Alerts > Data Ingestion and Replication Alerts** page in Operational Insights. Operational Insights then sends Data Ingestion and Replication alert notifications to the users and user groups you select, whenever an ingestion and replication job acquires one of the configured statuses, detects a DDL change, or a delay in the change data capture processing occurs.

Database ingestion and replication job details

You can drill down on a job to see job details from the **My Jobs** page accessed from the Data Integration Home page, from the **All Jobs** and **Running Jobs** pages in Monitor, and from the **Overview** and **All Jobs** tabs on the **Data Ingestion and Replication** page in Operational Insights.

After you click a job name, job information is displayed in the **Task Summary**, **Overview**, **Object Detail**, and **Alerts** panes. The **Alerts** pane is displayed only for incremental load and combined initial and incremental load jobs in Operational Insights.

The following image displays the panes:

Note: Clicking the expander arrow next to the pane name expands or collapses the **Task Summary**, **Object Detail**, and **Alerts** panes.

Task Summary

View a diagram that shows the source and target connector types and connection names. It also shows the calculated data throughput, in rows per second if the job has successfully replicated data to the target, regardless of the job's current status. If the calculated value is 0, indicating no data has flowed to the target, the throughput is not displayed.

In the *upper right corner*, the **Status** field displays the job status, which can be one of the following values:

- **Up and Running.** The job is running.
- **Running with Warning.** The job is running with a warning. This state can also occur when one or more table-specific subtasks fail but some subtasks are still running.
- **On Hold.** The job is in a paused state while the Database Ingestion and Replication (DBMI) agent is being updated.
- **Stopping.** The job is stopping in response to a Stop request.
- **Stopped.** The job was intentionally stopped.
- **Failed.** The job ended abnormally, the task deployment to the job failed, or one or more table-specific subtasks failed. Also, for an initial load job, the job was stopped.
- **Deploying.** The job is being deployed.
- **Deployed.** The job has been deployed.
- **Aborting.** The job is stopping immediately in response to an Abort request.
- **Aborted.** The job has been aborted.
- **Undeploying.** The job is being undeployed.
- **Undeployed.** The job has been undeployed.
- **Completed.** The job completed successfully.

Also, for incremental load jobs and combined initial and incremental load jobs, you can download the job execution log for the entire job run. In the *Select log* list under the diagram, select one of the following log types:

- **Complete Log.** The entire log, including all types of messages. It is available for any job that ran, regardless of its state.
- **Error.** The error log, which includes messages only for errors that occurred. It is available for Failed jobs only. Use this log to determine the reason for the job failure, for example, the deployment failed. If the log file ends with an ellipsis (...), the log has been truncated because of its long length. In this case, download the Complete Log to see all error messages.

Overview

View information about the job, including the runtime environment, task name, load type, task type, task definition location, user who started the job, records read, records written, number of subtasks, and schedule name, if a schedule was used.

The circle image displays the number of subtasks on source tables by status. The color of the circle's rim corresponds to the status.

The following table describes the summary job properties and statistics:

| Property | Description |
|----------------------|--|
| Duration (on left) | The amount of time, in the hh:mm:ss format, that the job ran before it ended. |
| Start Time (on left) | The date and time when the job was deployed. |
| End Time (on left) | The date and time when the job ended because it completed processing, was stopped, or failed. This field is not displayed for running jobs |
| Runtime Environment | The name of the runtime environment that the job uses to run. |
| Agent Name | For incremental load and combined initial and incremental load jobs, the name of the Secure Agent that the job uses to run. In case of a failover from one Secure Agent to another, Database Ingestion and Replication displays the latest Secure Agent assigned to the job. Note: For initial load jobs, the Agent Name is displayed in the Object Detail pane. For combined initial and incremental load jobs, the Agent Name is displayed both in the Task Summary and the Object Detail panes. |
| Task Name | The name of the associated database ingestion and replication. You can click the task-name link to view or edit task details in Data Ingestion and Replication, if necessary. If you edit the task, you must redeploy it for the updated task definition to be used for a job. |
| Load Type | The type of load operation that the job performs. Options are: <ul style="list-style-type: none"> - Initial Load. Loads a snapshot of source data read at a specific point-in-time to a target. - Incremental Load. Loads incremental data changes to a target on a continuous basis, until the job is stopped or ends. - Initial and Incremental Load. Performs an initial load and then automatically switches to an incremental load. |
| Task Type | The type of task, which is Database Ingestion and ReplicationTask . |
| Task Location | The project or project folder that contains the database ingestion and replication task definition. |
| Started By | The name of the user who started the job. |
| Records Read | The number of records that were read from the source. Note: For a combined initial and incremental load job, the Records Read count might be greater than the total number of object-level DML change records written. This behavior occurs because the initial load or resync processing always starts after change data capture has begun. As a result, some change records are included in the Records Read count and then discarded before initial load or resync processing starts. These discarded records cause the Records Write count to be less than the Records Read count. |
| Records Written | The number of records that were successfully replicated to the target. Note: The Records Written value might be different from the Records Read value if source records are discarded. For example, in a combined initial and incremental load job, change records read from the source before the initial unload phase completes are discarded because they're not yet needed. |
| Subtasks | The number of subtasks that the database ingestion and replication job used to replicate data from source tables to the target. When a job runs, it uses a separate subtask to process each source table. |

| Property | Description |
|------------------|---|
| Capture Progress | For incremental load and combined initial and increment load jobs, the date and time in the source change stream to which capture processing has progressed, as shown in the time zone of the user profile. |
| Schedule | For initial load jobs, the name of the schedule that is used to run the job or "No schedule" if you run the job manually. |

Object Detail

On the **Object Detail** pane lists subtasks on the source tables. You can view statistics and status information by source table from the last run of a database ingestion and replication job. When you click the expander arrow next to an object name, counts of processed inserts, updates, deletes, and LOB changes are shown for the table.

The following image shows a sample expanded object row:

| Object | Target Object | Status | Log |
|----------------------------------|--|---------|------------|
| DBMICDC_ppg013_rc_rongetype1_mri | KARTHEEK_ppgTC3_ppg013_rc_rongetype1_mri_uno=flake | Stopped | Select log |
| Inserts | Updates | Deletes | LOBs |
| 11 | 4 | 11 | 0 |
| DBMICDC_ppg014_rc_rongetype2_mri | KARTHEEK_ppgTC3_ppg014_rc_rongetype2_mri_uno=flake | Stopped | Select log |

Note: This pane shows information for the latest job run. This tab is blank for jobs that have not run or are resuming.

The following table describes the **Object Detail** fields that are displayed for each table, depending on the load type and status:

| Column | Description |
|---------------|---|
| Object | <p>The name of the source table or view for which data was propagated to the target.</p> <p>For an incremental load job or a combined initial and incremental load job, click the arrow icon to the left of the object name to display detailed counts of Inserts, Deletes, Updates, LOBs, and DDL statements processed. For a combined initial and incremental load job, the Unload Count field is also displayed to show the number of records that the initial load portion of processing read from the source. The following usage notes apply to the detailed CDC counts:</p> <ul style="list-style-type: none"> - The counts are only for the current job run. If you stop and restart the job, the counts start over from zero. Do not use these counts to identify the number of rows written to the target. - The counts are based on rows read from the source and do not reflect the records written to the target. Target write operations might be optimized by combining operations and reducing the number of physical writes. In this case, the counts might not match the number of write operations. - The value N/A means that the count value is not applicable for the count type or the value has not yet been calculated. - The Unload Count might not reflect the number of source records at the time the job is started or resynchronized because of a delay in the start of unload processing. Between the time of the unload request and start of unload processing, rows might be added to or deleted from the source table. |
| Target Object | The name of the target object that is mapped to the source object. |

| Column | Description |
|-----------------|--|
| Agent Name | <p>For initial load and combined initial and incremental load jobs, the name of the Secure Agent on which the job was run.</p> <p>For incremental load jobs, the Agent Name is displayed in the Overview section of the Task Summary pane. For combined initial and incremental load jobs, the Agent Name is displayed both in the Task Summary and the Object Detail panes.</p> |
| Records Read | <p>For an initial load job, the number of records that were read from the source. For other load types, this information is available only at the job-level on the Job Overview tab.</p> |
| Records Written | <p>For an initial load job, the number of records that were successfully written to the target. For other load types, this information is available only at the job-level on the Job Overview tab.</p> <p>Note: The Records Read value can be greater than the Records Written value if some records read from the source were discarded. For example, in a combined initial and incremental job, any source change records read before the initial unload phase of the job has completed are discarded.</p> |
| Task Duration | <p>For an initial load job, the amount of time the subtask that processed the source table ran before it completed or was stopped. For other load types, this information is available only at the job-level on the Job Overview tab.</p> <p>When a job runs, it uses a separate subtask to process each source table.</p> |
| Stage | <p>For a combined initial and incremental load job, this column shows the stage in the transition from initial load processing to CDC processing for the table-specific job subtask. This column does not appear for other load types.</p> <p>The stage can be one of the following values:</p> <ul style="list-style-type: none"> - Not Started. Initial load processing has not yet started for the table, or if an error occurred and the table is in the Error on Retry state, the next attempt to process the table has not yet started. - Started. Initial load processing has started. - Unloading. The subtask is unloading data from the table as part of initial load processing. - Unloaded. The subtask has finished unloading data from the table as part of initial load processing. - Completed. The subtask completed initial load processing of the table. - Normal. The subtask completed initial load processing of the table and has started CDC processing of the table. - Cancelled. Initial load processing was cancelled or stopped. - Error. The subtask detected an error in the source table. <p>Actions menu > Resync</p> <p>For a subtask in a combined initial and incremental load job, if the subtask stage is Normal and the subtask status is any status other than Queued or Starting, the Actions menu is displayed on the right end of the subtask row. From the Actions menu, you can select Resync to resynchronize the source and target objects. For more information, see "Resynchronizing source and target objects" in Data Ingestion and Replication help.</p> |

| Column | Description |
|-----------------|--|
| Status | <p>The status of the job subtask for the source object.</p> <p>Note: If the job stops running, the subtask status reflects the status last collected before the job ended. For example, the job might be aborted but the subtask is in a Running status.</p> <p>The state can be one of the following values:</p> <ul style="list-style-type: none"> - Queued. The subtask has not yet started running. - Starting. The subtask is starting. - Started. For a combined initial and incremental load job, the subtask has started. - Running. The subtask is running. - On Hold. The subtask, as well as the job, is in a paused state while the Database Ingestion and Replication (DBMI) agent is being updated. - Completed. The subtask completed processing successfully. - Stopping. The subtask is stopping in response to a Stop request. - Stopped. The subtask has stopped. - Aborting. The subtask is ending immediately in response to an Abort request. - Aborted. The subtask has been aborted. - Failed. The subtask ended unexpectedly. - Error. The subtask is in error and no longer writing data to the target table. For a combined initial and incremental load job, the subtask might be running and processing incremental change data but no data is being sent to the target. - Error on Retry. An error occurred on the last retry of subtask processing, and now the subtask is waiting to retry processing again. <p>Note: If a DDL change occurs on a source table and then you resume the job, the table subtask state might not change as expected until the first DML operation occurs on the source table.</p> |
| Data Validation | <p>For initial load jobs that have an Oracle or a SQL Server source and a Snowflake target, if the task status is Completed, you can run data validation to compare the source and target data. This column shows the status of the data validation task. This column does not appear for other load types and other source and target combinations, or if the task is in any other status than Completed.</p> <p>From the Actions menu displayed on the right end of the subtask row, you can select Run Data Validation to start the validation task. For more information, see "Running data validation for a database ingestion and replication job" in Data Ingestion and Replication help.</p> <p>The data validation status can be one of the following values:</p> <ul style="list-style-type: none"> - N/A. The data validation has not been run. - Processing. The data validation is running. - Success. The data validation completed successfully. Click on the Success status to see the data validation summary. To download the data validation report, click the Download icon. - Error. An error occurred during the data validation execution. To view the error message, click the Download icon. - Report Pending. The data validation has been requested but the Data Validation service on the Secure Agent that runs the task is down. - No Comparable Data. The source or target table does not contain any data. |


| Column | Description |
|--------------|---|
| CDV Last Run | <p>If data validation is available for the task, this column shows the date and time when the data validation was last performed for the given subtask. If the data validation has not been run, the field value is N/A.</p> <p>This column appears in initial load jobs that have an Oracle or a SQL Server source and a Snowflake target, only if the task status is Completed.</p> |
| Log | <p>You can download a job execution log for a source object. The type and availability of the log depends on the load type and status. Options are:</p> <ul style="list-style-type: none"> - Complete. The complete log for an object subtask from job execution. This log type is available for a Completed, Failed, or Aborted subtask in an initial load job. - Error. The log that contains error messages. This log type is available only for a Failed or Error subtask in an initial load or incremental load job. - Stage Log. The log that covers the transition from the initial processing phase to the incremental processing phase in a combined initial and incremental load job for a source object. <p>To download a log, click the Download icon.</p> <p>Note: If you undeployed the job, you can download the log for a table only if the associated task has not been deleted.</p> <p>For incremental load jobs, you can get the complete log and error log for the entire job run from the Task Summary pane.</p> |

Note: To control the line spacing in the list, click the **Settings** icon next to the **Find** box.

Alerts

The **Alerts** pane appears on the **Data Ingestion and Replication** page in Operational Insights for the selected incremental load or combined initial and incremental load job. On the **Alerts** pane, you can view alert messages that appear for certain events, such as source schema changes, during incremental load or combined initial and incremental load processing.

The following image shows this pane with an alert example:



Note: The **Alerts** pane displays alert messages for all detected schema changes even if you set the schema drift options for the associated task to Ignore.

You can filter the list of alerts based on severity or a date range. To specify a date range, enter one of the following types of values in the **Filter** field:

- **Any Time** for all stored alerts.
- **Today** for alerts issued today from midnight to 11:59 pm.
- **Last Week, Last Month, or Last Year** to show alerts from the beginning of last week, month, or year to present.
- **Custom** to specify a custom date range that consists of a beginning date and time and an ending date and time.

The list of alerts includes the following columns:

| Column | Description |
|--------|---|
| Level | Severity level of the alert message, such as Critical or Warning. |
| Code | Alphanumeric code that identifies the alert type followed by the date and time when the event occurred. |

Click the expander arrow to display a description of the event.

You can configure alert notifications for database ingestion and replication jobs from the **Alerts > Data Ingestion and Replication Alerts** page in Operational Insights. Operational Insights then sends Data Ingestion and Replication alert notifications to the users and user groups you select, whenever a database ingestion and replication job acquires one of the configured statuses, detects a DDL change, or a delay in the change data capture processing occurs.

File ingestion and replication job details

You can view job details and results for each file ingestion and replication job instance, including the status of the job and success and error statistics.

To view detailed information about a file ingestion and replication task, click the task name on the **My Jobs** page that's accessed from the navigation bar in Data Integration Home page, the **All Jobs** and **Running Jobs** pages in Monitor, or from either tab of the **Data Ingestion and Replication** page in Operational Insights. You can also download the session log.

The following image shows the details of a file ingestion and replication job in Operational Insights:

The screenshot displays the details of a file ingestion and replication job. The job name is `adlsgen2_to_ftps_1SatJun122021094156-124505`. The job type is `File Ingestion Task`, started by `filelistener` on `Jun 12, 2021, 10:31:16 AM`. It ended on `Jun 12, 2021, 10:31:31 AM` with a duration of `00:00:15`. The job status is `Success`. The results show `14` success files, `0` error files, and `0` duplicate files. The error message is `Job completed normally`. Below the job details is a table of file events:

| Name | File Size(Bytes) | Status | Transfer Type | Start Time | Duration (ms) | Remarks |
|-----------|------------------|---------|---------------|---------------------------|---------------|---------|
| File7.txt | 26 | Success | FTPS Upload | Jun 12, 2021, 10:31:30 AM | 358 | |
| File6.txt | 32 | Success | FTPS Upload | Jun 12, 2021, 10:31:30 AM | 265 | |
| File7.txt | 26 | Success | ADLS Download | Jun 12, 2021, 10:31:29 AM | 929 | |

Job Properties

The job properties for a file ingestion and replication job instance include general information about the job, when it started and ended, and its duration.

The following table describes the job properties:

| Property | Description |
|------------|--|
| Task Name | The name of the associated file ingestion and replication task. You can click the task-name link to view or edit task details in Data Ingestion and Replication. |
| Task Type | Task type. In this case, file ingestion and replication task. |
| Started By | Name of the user or schedule that started the job. |
| Start Time | Date and time when the job was started. |
| End Time | Date and time when the job completed or stopped. |
| Duration | The amount of time the job ran before it completed or was stopped. |

Results

The job results for a file ingestion and replication job instance include the job status and error statistics.

The following table describes the job results:

| Property | Description |
|-----------------|---|
| State | <p>Job status. A job can have one of the following statuses:</p> <ul style="list-style-type: none">- Running. The job is still running.- Success. The job completed successfully.- Failed. The job did not complete because it encountered errors- Aborted. The job was aborted. <p>Note: When the Secure Agent is unavailable while running a file ingestion job, the job remains in an unresponsive state, and after 200 minutes, its status changes to Aborted.</p> |
| Session Log | <p>Allows you to download the session log file. By default, Informatica Intelligent Cloud Services stores session logs for 10 runs before it overwrites the logs with the latest runs. If you need the session logs for earlier runs, take a backup of the directory that holds the session log files.</p> <p>Session log files are written to the following directory: <Secure Agent installation directory>/apps/MassIngestionRuntime/data/taskLogs</p> |
| Success Files | Number of files that are successfully transferred, downloaded, and uploaded to the target. |
| Error Files | Number of files that were not transferred to the target. |
| Duplicate Files | Number of files that were identified as duplicates. |
| Error Message | Error message, if any, that is associated with the job. |

File Events

This section shows the total number of files that the file ingestion and replication job transferred with information about each file.

The File Events section is updated each time the file ingestion and replication job transfers a file. The state of the file is updated throughout the file transfer process. You can track the progress of a file transfer based on the state of the file.

The File Events section displays the following properties for each file:

| Property | Description |
|---------------|--|
| Name | The name of the file. |
| File size | The size of the file in bytes. |
| Status | <p>The status of the file transfer. A file can have one of the following status:</p> <ul style="list-style-type: none">- Success. The file transfer completed successfully.- Failed. The file transfer did not complete because it encountered errors.- Processing. The file transfer is still running.- Duplicate. The task previously transferred a file with the same name, directory location, and size.- Interrupted. The file transfer is interrupted because of network issues or changed server credentials during the file transfer. Run the file ingestion and replication job to resume the transfer of the interrupted files. <p>Note: The status is applicable when the file ingestion and replication job transfers file from or to the advanced FTP, advanced SFTP, or advanced FTPS sources and targets.</p> <ul style="list-style-type: none">- In Doubt. The previous task instance encountered errors while transferring the file. Applicable for tasks where the source is configured to skip duplicate files.- Quarantined. The job marks any infected file it detects from a source as quarantined. <p>You can monitor the Status property to track the progress of the file transfer of each file.</p> |
| Transfer Type | <p>The type of file transfer. A file can have one of the following transfer types:</p> <ul style="list-style-type: none">- <code><source_name></code>Download. The file is downloaded from source. <code><source_name></code> is the name of the source.- <code><target_name></code>Upload. The file is uploaded to the target. <code><target_name></code> is the name of the target.- Copy from Source. The file ingestion and replication task is performing file processing actions.- Copy to Target. The file is transferred from a local directory to a local directory. |
| Start time | Date and time when the file transfer started. |
| Duration | The length of time to transfer the file, in milliseconds. |
| Remarks | Applies to file events in Failed status. The message includes the reason for failure of the event based on the file transfer type. |

Note: If you click the instance name of a file ingestion and replication job on the **My Jobs** page, accessed from the Data Integration Home page, the job details page appears. On this page, you can click the name of the associated task to open the task definition in edit mode if you need to change the task and re-run the job.

Streaming ingestion and replication job details

You can drill down on a job to see job details from the **My Jobs** page accessed from the Data Integration Home page, from the **All Jobs** and **Running Jobs** pages in Monitor, and from the **Overview** and **All Jobs** tabs on the **Data Ingestion and Replication** page in Operational Insights.

After you click a job name, job information is displayed on the **Overview** tab with a summary source-to-target processing diagram at the top. Information is also displayed on the **Alert**, **Performance**, and **Past Run** tabs if available.

The following image shows the diagram and **Overview** tab:

Overview tab

The **Overview** tab displays general properties of the job. You can download the job log too.

The following table describes the job overview properties:

| Property | Description |
|---------------|---|
| Job Name | The name of the job. |
| Version | The version number of the job. |
| Task Type | The task type of streaming ingestion and replication task. |
| Task Location | The project or project folder that contains the streaming ingestion and replication task. |
| Started By | The name of the user who deployed the job. |
| Secure Agent | The location where the Secure Agent is running. A warning symbol near the Secure Agent indicates that the Secure Agent is either offline or not reachable. |

| Property | Description |
|---------------------|--|
| Status | The status of the job. A job can have one of the following statuses: <ul style="list-style-type: none"> - Deploying. The job is being deployed. - Up and Running. The job is running. - Running with Warning. The job is running with warnings. - Running with Error. The job is running with error. If a job continuously runs with warnings for seven minutes or for the time specified in the runtime option, the state of the job changes to Running with Error. - Undeployed. The job is undeployed. - Stopped. The job was intentionally stopped. |
| Duration | Total time the job ran before it is undeployed. The total time is shown in <code>hh:mm:ss</code> format. |
| Start Time | The date and time when the job was deployed. |
| Runtime Environment | Name of the runtime environment that the job uses to run. |
| Download Log | You can download one of the following types of logs: <ul style="list-style-type: none"> - Complete. The entire log, including all types of messages. It is available for any job that ran, regardless of its status. - Latest. Latest version of the log. To download a log to your local system, click the Download icon. |

Alert tab

The **Alert** tab displays the alert messages when an event occurs.

The following image shows the **Alert** tab for a streaming ingestion and replication job:




The following table describes the job alert properties:

| Property | Description |
|----------|--|
| Alert | The messages or a group of messages that the job returns when a deployed job encounters a warning. |
| Time | The date and time when the event occurred. |

Performance tab

The **Performance** tab displays graphs of throughput information for the source and target of the job.

The following image shows the **Performance** tab:



The following table describes the job performance properties:

| Property | Description |
|----------------|---|
| Total messages | The average number of messages streamed per second. |
| Total kb | The average kilobits of messages streamed per second. |

Past Run tab

The **Past Run** tab displays the statistics and status information for previous runs of a streaming ingestion and replication job.

The following table describes the properties displayed for each run:

| Column | Description |
|------------|--|
| Version | The version number of the job. |
| Start Time | The date and time when the job was deployed. |
| End Time | The date and time when the job was undeployed. |
| Duration | Total time the job ran before it is undeployed. The total time is shown in <code>hh:mm:ss</code> format. |

Data Ingestion and Replication alerts

You can configure Operational Insights to send Data Ingestion and Replication alert notifications to users when application ingestion and replication jobs or database ingestion and replication jobs detect certain status changes, DDL changes, or delays in processing DML data. For example, you can alert an administrator when the job status changes to Failed or Running with Warning.

For you to configure alerts, your organization must have the OperationalInsightsAdvancedEdition license with the Job Alerts feature. Also, to create or modify alerts, you must have the one of the following roles:

- Admin role
- A custom role that has the Operator role along with the Read privilege for the User and Group assets for Administrator service.

These role requirements do not pertain to users who only read alerts.

From the **Data Ingestion and Replication Alerts** tab on the **Alerts** page, you can create, edit, or delete alert rules. To create an alert rule, click **Create Alert**. To edit or delete a listed alert rule, click the pencil or trash can icon at the right end of the rule row in the list.

You can configure an alert rule to send alerts when one or more of the following types of events occur:

- The status of an application ingestion and replication job or database ingestion and replication job changes to a status that you selected for alerting.
- A DDL schema change event occurs for incremental load or combined initial and incremental load jobs that have schema drift enabled.
- An incremental load or a combined initial and incremental load job has not been processing any DML data for a significant period of time.

Also specify the recipients of the alert notifications and whether the alert rule applies to the entire organization or to a selected ingestion task.

Operational Insights polls for application ingestion and replication and database ingestion and replication alerts every 5 minutes.

Configuring alerts for Data Ingestion and Replication jobs

You can configure alerts for application ingestion and replication jobs and database ingestion and replication jobs to notify users about the job status, DDL changes, and delays in CDC processing. Alerting is not supported for file ingestion and replication jobs and streaming ingestion and replication jobs.

1. On the **Alerts** page, click the **Data Ingestion and Replication Alerts** tab.
2. Click **Create Alert**.
3. On the **Create Alert Rule** page, configure the alert details:
 - a. In the **Alert Rule Name** field, enter a name for the alert. Maximum length is 255 characters.
 - b. In the **Alert Rule Description** field, enter an optional description for the alert. Maximum length is 255 characters.
 - c. In the **Ingestion Type** field, select **Application Ingestion and Replication Task** or **Database Ingestion and Replication Task** or both options.
 - d. In the **Rule Applies To** field, configure the alert scope by selecting one of the following options:
 - **Entire Org**. Apply the alert rule to all jobs in the organization.
 - **Task Asset**. Apply the alert rule to the jobs associated with the ingestion task that you select.
4. In the **Alert Enabled** field, enable or disable the alert rule. Disable the alert if you do not want alerts to start being sent right away. By default, the alert is enabled.
5. Configure alert conditions in one or more of the following ways:
 - To alert based on a job status change, select each job status for which to send an alert. By default, only **Failed** is selected.
 - To alert when a source schema change is detected, select **DDL**. This option applies to incremental load and combined initial and incremental load jobs for which schema drift options are enabled.
 - To alert when an incremental load or a combined initial and incremental load job has not been processing any DML data for a significant period of time, select **Source CDC Delay**.
6. In the **Send Email To** field, select one or more users or user groups to send email notifications for the configured alerts to.

You can select multiple users individually, one or more user groups, or any combination of individual users and user groups. To be available for selection, the users and user groups must have been previously defined in Administrator service.

CHAPTER 8

Monitor MDM SaaS

You can use Operational Insights to view data processing analytics and monitoring statistics for MDM SaaS.

You can use data processing analytics to assess the effectiveness of your investment in Informatica by analyzing usage. You can use monitoring statistics to assess the health of your organization and troubleshoot job failures. You can monitor various usage statistics such as the number of source records, master records, and users.

Prerequisites to monitor MDM SaaS

The prerequisites to monitor MDM SaaS usage and jobs vary based on how you monitor MDM SaaS. To access Operational Insights to monitor MDM SaaS usage and jobs, your user role requires necessary permissions. To monitor MDM SaaS in a business application as a predefined dashboard, you don't require any additional permissions.

To provide monitoring access to a custom user role, perform the following steps:

1. In Administrator, click **User Roles**.
2. Select the custom user role for which you want to provide monitoring access.
The role details page appears.
3. From the **Services** list, select **Operational Insights**.
4. Click the **Features** tab.
5. Select **Operational Insights - view**.
6. From the **Services** list, select **MDM Configuration**.
7. Select **Configuration and Authoring**.
8. Click **Save**.

Monitor usage statistics for MDM SaaS

You can view a usage summary that shows the counts of source records, master records, and users. You can also view the usage details consisting of source and master record counts grouped by business entity.

The **Overview** tab shows an overview of MDM SaaS usage statistics.

The following table describes the usage statistics that you can see in the **Usage Summary** panel of the **Overview** tab:

| Usage Statistic | Description |
|-----------------|---|
| Source Records | The number of source records. The source records contribute to the master records. The number doesn't include the records that the ingress job rejects. |
| Master Records | The number of master records. A master record is the consolidated version of source records from multiple external systems. |
| Users | The number of users in the organization that you are monitoring. |

The **Usage Details** panel on the **Overview** tab shows the source and master record counts by business entity.

Monitor MDM SaaS jobs

When you monitor MDM SaaS jobs, you can view a summary of jobs, and view all job instances and job schedules. You can also view key metrics for each job type.

You can monitor the data that each job processes or compare the metrics of various jobs. For example, you can compare the number of records processed by two or more ingress jobs to determine the number of records that were imported from source systems.

View job summaries

You can view the number of jobs that ran successfully, have failed, are running, and are scheduled to run in the next 24 hours on the job summary page. You can also view the top 5 jobs based on their run duration and the number of records processed.

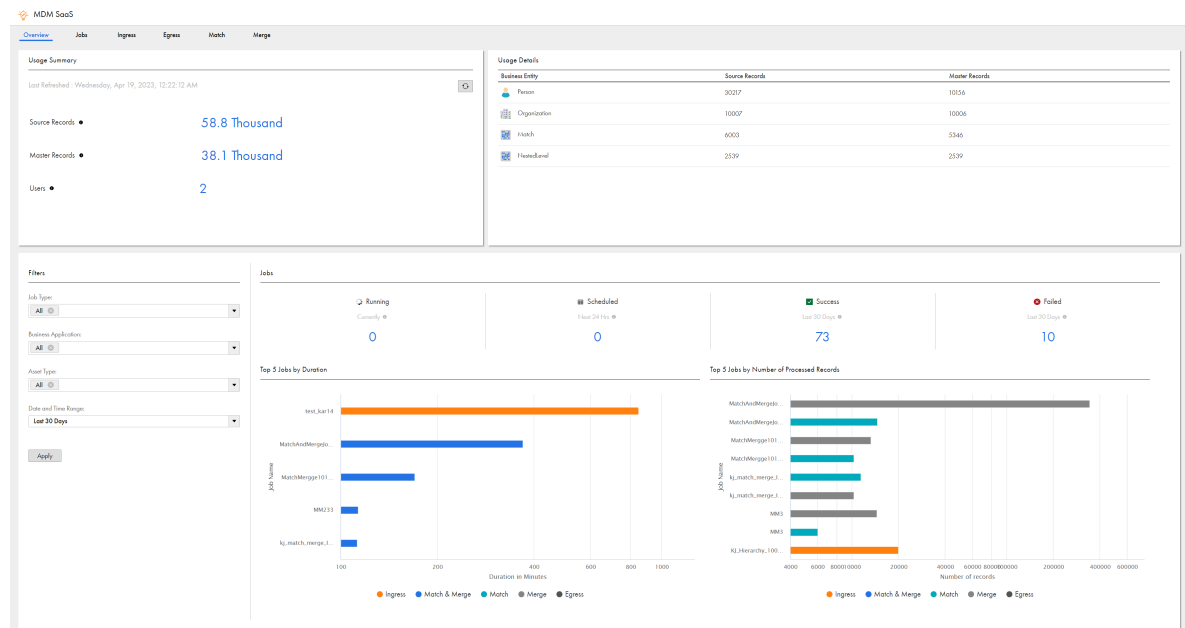
To view the summary of jobs, use the **Overview** tab of the **MDM SaaS** page. You can filter and view the job summaries.

The following table describes the job filters:

| Filter | Description |
|----------------------|--|
| Job Type | Filters the job details based on the selected job type. |
| Business Application | Filters the job details based on the selected business application. The values depend on the business applications that your organization subscribes to. |

| Filter | Description |
|---------------------|---|
| Asset Type | Filters the job details based on the selected asset type. |
| Date and Time Range | <p>Filters the following job details based on the selected date and time range:</p> <ul style="list-style-type: none"> - The number of successful and failed jobs. - The top 5 jobs based on their run duration and the number of records processed. <p>You can choose to set a custom date and time range.</p> <p>The specified data and time range don't affect the number of running and scheduled jobs. The job summary page displays the number of running jobs and the number of jobs scheduled to run in the next 24 hours, irrespective of the applied filters.</p> |

The following image shows the **Overview** tab:



View job instances and schedules

You can view a list of job instances and job schedules on the **Jobs** tab.

When you view the list of job instances and job schedules, you can monitor the status of jobs and which jobs have schedules.

The following image shows the **Jobs** tab:

MDM SaaS

Overview Jobs Ingress Egress Match Merge

Job Schedules

Job Instances (556) Updated On Apr 14, 2023, 10:49:53 AM

| Name | Job Type | Started By | Start Time | End Time | Status |
|---|-----------------------|---------------------|--------------------------|--------------------------|---------|
| Report internal masterRecords_count_by_BE | Report generateReport | GladionRaiSystemPod | Apr 14, 2023, 7:25:51 AM | Apr 14, 2023, 7:25:56 AM | Success |
| Report internal sourceRecords_count_by_BE | Report generateReport | GladionRaiSystemPod | Apr 14, 2023, 7:25:41 AM | Apr 14, 2023, 7:25:45 AM | Success |
| MM233 | Match and Merge | GladionRaiSystemPod | Apr 14, 2023, 3:41:11 AM | Apr 14, 2023, 5:34:43 AM | Success |
| MatchMergeOn14Apr | Match and Merge | GladionRaiSystemPod | Apr 14, 2023, 2:03:04 AM | Apr 14, 2023, 2:03:04 AM | Failed |
| Hierarchy Import file.csv | Ingress | GladionRaiSystemPod | Apr 14, 2023, 1:39:21 AM | Apr 14, 2023, 1:39:53 AM | Success |
| MatchMergeOn14Apr | Match and Merge | GladionRaiSystemPod | Apr 14, 2023, 3:14:28 AM | Apr 14, 2023, 3:14:03 AM | Success |
| Report internal masterRecords_count_by_BE | Report generateReport | GladionRaiSystemPod | Apr 13, 2023, 4:10:02 PM | Apr 13, 2023, 4:10:07 PM | Success |
| Report internal sourceRecords_count_by_BE | Report generateReport | GladionRaiSystemPod | Apr 13, 2023, 4:09:54 PM | Apr 13, 2023, 4:09:58 PM | Success |
| Hierarchy Import file.csv | Ingress | GladionRaiSystemPod | Apr 13, 2023, 1:40:40 PM | Apr 13, 2023, 1:41:10 PM | Success |
| MatchMergeJob | Match and Merge | GladionRaiSystemPod | Apr 13, 2023, 3:31:25 PM | Apr 13, 2023, 3:17:06 PM | Success |

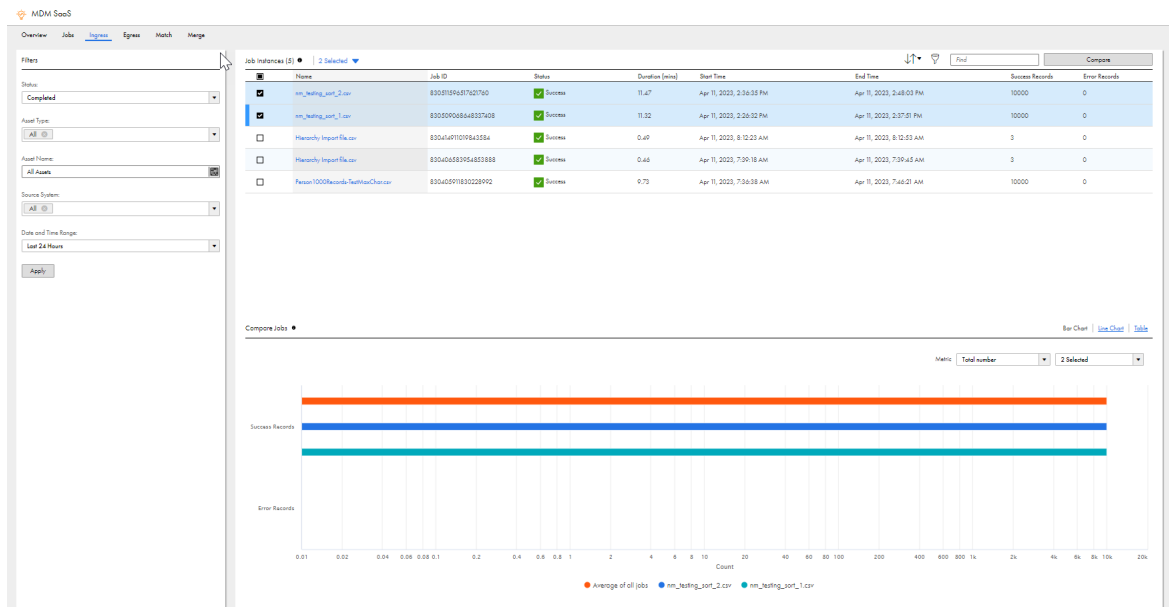
1 - 25 of 556 Items Page 1 of 23 Items Per Page: 25

You can filter the list of jobs by job type or status. You can filter the list of job schedules by schedule name, job name, or status. To specify a filter, click **Filter**, then click **Add Filter** and select the property to filter by.

View key job metrics

You can view key job metrics, such as the number of records processed, for ingress, egress, match, and merge jobs. To view job metrics, use the **Ingress**, **Egress**, **Match**, and **Merge** tabs.

The following image shows the **Ingress** tab with key metrics of ingress jobs:



You can view the status, start time, end time, duration, and metrics of a job. The metrics are specific to each job type. You can also sort and filter the jobs.

The following table describes the job filters:

| Filter | Description |
|---------------------|---|
| Status | Filters the job details based on the selected job status. |
| Asset Type | Filters the job details based on the selected asset type. |
| Asset Name | Filters the job details based on the selected asset name. The values depend on the assets in your organization. |
| Source System | Filters the job details based on the selected source system for an ingress job. The values depend on the source systems from which your organization gets data. |
| Date and Time Range | Applicable to completed jobs. Filters the job details based on the selected date and time range. You can choose to set a custom date and time range. |

You can also filter the jobs by job ID, status, start time, end time, duration, and metrics of a job. To specify a filter, click **Filter**, then click **Add Filter** and select the property to filter by.

Ingress job metrics

You can analyze the ingress job metrics to determine the number of records that were imported from source systems.

The following table describes the metrics that you can view for an ingress job:

| Metric Name | Description |
|-----------------|--|
| Total Records | The total number of records that were processed. |
| Success Records | The number of records that were successfully imported. |
| Error Records | The number of records that failed to import. |

Egress job metrics

You can analyze the egress job metrics to determine the number of records that were exported from the Business 360 data store to other external data sources.

The following table describes the metrics that you can view for an egress job:

| Metric Name | Description |
|----------------------------|--|
| Number of Exported Records | The number of records that were successfully exported. |

Match job metrics

You can analyze the match job metrics and decide whether to modify the match model configuration.

The following table describes the metrics that you can view for a match job:

| Metric Name | Description |
|---|---|
| Total Matched Records | The number of records that were successfully matched. |
| Failed Matched Records | The number of records that failed to match. |
| Record Pairs for Manual Merge | The number of record pairs that require manual review. |
| Record Pairs for Automated Merge | The number of record pairs that will automatically merge based on declarative rules. |
| Record Pairs for Machine Learning Model Match | The number of record pairs that will automatically merge based on the machine learning model. |

Merge job metrics

You can analyze the merge job metrics to determine the number of record pairs and record pair groups that were created by the job.

The following table describes the metrics that you can view for a merge job:


| Metric Name | Description |
|--------------------|---|
| Total Record Pairs | The number of record pairs that were successfully created by the job. |
| Number of Groups | The number of record pair groups that were successfully created by the job. |

View job performance

You can view the performance of a job in the **Job Instances** panel. You can view the number of records, record pairs, or groups processed in a time interval of a job.

To view the performance of a job, in the **Job Instances** panel, hover over a job, then click **More Actions > View Job Performance**. You can also view the performance of a job by clicking the job duration link in the **Duration (mins)** column.

The following image shows the **Job Performance** window:



Export jobs

You can export the list of jobs and the details of those jobs in the **Job Instances** panel to a CSV file.

The following table describes the export types and the corresponding task steps:

| Export Type | Task |
|--------------------|--|
| All jobs | Click None Selected > Select All , then click All Selected > Export . |
| All jobs on a page | Click None Selected > Select All On This Page , then click All Selected > Export . |
| Selected jobs | Click [Number of jobs selected] Selected > Export . |


Compare jobs

You can compare two to five jobs and view their key metrics in a bar chart, line chart, or table. The key metrics vary based on the type of jobs.

For example, you can compare the success records between two ingress jobs or the number of record pairs that require manual review between two match jobs.

To compare jobs, select two to five jobs in the **Job Instances** panel, then click **Compare**.

The following image shows the **Compare Jobs** panel:



Part III: Monitoring on-premises applications

This part contains the following chapters:

- [Register and manage domains, 108](#)
- [Monitor Data Engineering Integration domains , 121](#)
- [Monitor Data Quality domains , 124](#)
- [Monitor PowerCenter domains, 127](#)

CHAPTER 9

Register and manage domains

You must register each Informatica domain that you want to monitor using Operational Insights. A wizard guides you through the process. Domain registration is required only for on-premises products. You do not register domains for Informatica Intelligent Cloud Services.

The following are the key steps in the domain registration process:

1. Configure the connection to the domain.
2. Supply details that enable users to more easily locate the domain.
3. Configure the collectors.

After you complete the registration process and click **Save** as the final step, the collectors begin collecting and uploading data to Informatica Intelligent Cloud Services for use by Operational Insights.

Note: If you upgrade the Informatica domain from earlier versions to 10.4.0, you must edit the domain registered in Operational Insights to reflect the details of the upgraded Informatica domain version. To collect statistics for a Data Engineering Integration version 10.2.2 domain, you must apply EBF-14386 to each Data Engineering Integration node in the domain.

Enable the monitoring Model Repository Service

You must configure the monitoring Model Repository Service and the associated Model Repository in a domain before you register a domain with Operational Insights.

Operational Insights extracts CPU and memory consumption metrics for domain nodes from the Model repository specified in the Monitoring Configuration for the domain. The Model Repository Service that manages with the Model repository is known as the monitoring Model Repository Service.

You can create a monitoring Model Repository Service when you run the installer to create a domain. For more information, see the "Prepare for Application Services and Databases" chapter in the *Informatica Installation for Data Engineering Guide*.

You can also use the Administrator tool to configure the monitoring Model Repository Service in a domain. For more information, see the "Configuring the Monitoring Model Repository Service" section in the *Informatica Application Service Guide*.

Configuring the domain connection

Enter the information required to enable Operational Insights to connect to the domain. Operational Insights must be able to connect to the domain as an Informatica administrator.

1. Click **All Infrastructure** on the Operational Insights navigation bar.
2. Click **Domains > Register Domain**.
3. Enter the following general properties:

| Property | Description |
|---------------------|--|
| Domain Display Name | Domain name to display in the Operational Insights user interface. You can assign any name you like, but it must be unique within Operational Insights. |
| Domain Name | Domain name displayed in Informatica Administrator (the Administrator tool). |
| Master Gateway Host | Host name of the master gateway node machine. Enter the value exactly as shown in the General Properties > Host Name property for the node in the Administrator tool. To find this value in the Administrator tool: <ul style="list-style-type: none">- Select the Services and Nodes view, and then select the node in the Domain Navigator.- Under General Properties, locate the Host Name property. |
| Gateway Node Port | HTTP port used by the gateway node. |
| Domain Version | Informatica release installed in the domain. Operational Insights can monitor assets within all Informatica release 10.x domains. |
| Products | Informatica products to monitor using Operational Insights. Products are selected based on the domain version. |

4. Select the Secure Agent that collects and uploads data from the domain to the Informatica Intelligent Cloud Services.

| Property | Description |
|--------------------|--|
| Secure Agent Group | Group the Secure Agent installed in the domain belongs to. |
| Secure Agent Name | Name of the Secure Agent installed in the domain. |

5. Enter the following domain security details.

| Property | Description |
|-------------------------|---|
| Security Domain | Select the security domain used by the domain. |
| Administrator User Name | User name for the Informatica domain administrator account. |
| Administrator Password | Password for the Informatica domain administrator account. |

| Property | Description |
|---------------------|---|
| TLS Enabled | Select if the domain is secured with the Transport Layer Security (TLS) protocol |
| Truststore Path | If the domain is secured with TLS, copy the <code>infa_truststore.jks</code> file from a domain node to the Secure Agent host, and then specify the path and file name for the file on the Secure Agent host. By default, the file is installed in the following directory on each domain node: <Informatica installation directory>\services\shared\security |
| Truststore Password | If the domain uses a custom truststore file, specify the encrypted truststore password. |

6. Enter the following domain auto-scaling details to enable elastic nodes added to the grid to communicate with the domain.

| Property | Description |
|-------------------|---|
| Sitekey Path | Copy the <code>sitekey</code> file from a domain node to the Secure Agent host, and then specify the path and file name for the file on the Secure Agent host. By default, the file is installed in the following directory on each domain node: <Informatica installation directory>\isp\config\keys |
| Keystore Path | If the domain is secured with TLS, copy the <code>infa_keystore.jks</code> file from a domain node to the Secure Agent host, and then specify the path and file name for the file on the Secure Agent host. By default, the file is installed in the following directory on each domain node: <Informatica installation directory>\services\shared\security |
| Keystore Password | If the domain uses a custom keystore file, specify the encrypted keystore password. |

7. Click **Test Connection** to test the connection to the master gateway node.

Entering the domain details

Enter details to help users find the domain within Operational Insights.

You can use tags to categorize domains. Users can use the tags to search for the domains. You also select the domain type to indicate how the domain is used.

You can also organize domains by geographical location on an interactive map. The map is displayed on the Operational Insights home page. Assigning domains to locations helps you analyze performance and determine capacity and processing capabilities across the enterprise.

1. Assign tags to the domain under Domain Details. You can assign multiple tags to a domain.
 - To assign an existing tag, select it from the list.
 - To add a new tag, type the tag in the entry field, then click the **Enter** key on your keyboard.
2. Select the domain type that best matches how the domain is used within the organization..

3. Specify the location of the domain on the map.
 - To assign the domain to an existing location, click the location on the map.
 - To assign the domain to a new location, click where you want to add the location on the map, then type in the location name.
4. Click **Next** to save your entries.

Configure the Domain Configuration Collector

Configure the Domain Configuration Collector, which collects and uploads configuration metadata for the domain and all domain assets.

The default collection frequency is every 24 hours. You can create a custom schedule if needed to better suit your requirements.

The collector is enabled by default. You cannot disable the collector.

Configure the Collector Schedule

You can configure a custom schedule for the collector. The schedule you create overrides the default collector schedule.

Enter the following properties:

| Property | Description |
|-------------------|---|
| Repeats | The interval at which to repeat collection. |
| Repeats Frequency | The frequency at which to perform collection. The frequency is based on the repetition value you select. For example, to collect data every two hours, select Hourly as the repetition, and then set the frequency value to 2. |
| Starts on | The date and time the custom schedule takes effect. |
| Timezone | The timezone the schedule is based on. |

Configure the Domain Health Statistics Collector

Configure the Domain Health Statistics Collector, which collects and uploads availability statistics for domain assets.

The default collection frequency is every 5 minutes. You can create a custom schedule if needed to better suit your requirements.

The collector is enabled by default. Clear the **Enabled** checkbox to diable the collector.

Configure the Collector Schedule

You can configure a custom schedule for the collector. The schedule you create overrides the default collector schedule.

Enter the following properties:

| Property | Description |
|-------------------|---|
| Repeats | The interval at which to repeat collection. |
| Repeats Frequency | The frequency at which to perform collection. The frequency is based on the repetition value you select. For example, to collect data every two hours, select Hourly as the repetition, and then set the frequency value to 2. |
| Starts on | The date and time the custom schedule takes effect. |
| Timezone | The timezone the schedule is based on. |

Configure the Domain Resource Usage Statistics Collector

Configure the Domain Resource Usage Statistics Collector, which collects and uploads CPU and memory consumption metrics for all nodes within the domain.

The metrics are extracted from the Model repository specified in the Monitoring Configuration for the domain. You must configure the connection to the Model repository.

The default collection frequency is every 1 hour. You can create a custom schedule if needed to better suit your requirements.

The domain is enabled by default. Clear the **Enabled** checkbox to disable the collector.

Collecting historical data

You can configure the collector to populate Operational Insights with up to 60 days of historical data.

Historical data is collected for the previous 30 days by default. However you can specify any number of days between 1 and 60.

Data collection begins at the time the domain is added to Operational Insights. Roughly 24 hours worth of data is collected every hour, meaning that approximately 30 hours are required to populate Operational Insights with data for the prior month.

Historical data collection is enabled by default. Clear the **Collect Historical Data** checkbox to disable historical data collection.

Connecting to the Monitoring Statistics Model repository

Enter the information required to collect CPU and memory consumption metrics for domain nodes from the Model repository associated with the monitoring Model Repository.

You only need to enter the connection information when configuring either the Domain Resource Usage Statistics Collector or the Data Engineering Integration Collector. The same information is used by both collectors.

Use the Administrator tool that you use to manage the domain to locate the required property values:

- Select the **Services and Nodes** view.
- Click the **Monitoring Configuration** tab, and then note the name of the Model Repository Service.
- Select the Model Repository Service instance in the Domain Navigator, and then note the properties listed under Repository Database Properties.

After you locate the required property values, complete the following steps configure the connection to the Model repository.

1. Enter the following required properties:

| Property | Description |
|------------------------|--|
| Database Type | Model repository database type. |
| Username | User name for the Model repository database. |
| Password | Password for the Model repository database. |
| JDBC Connection String | <p>JDBC connection string used to connect to the Model repository database.</p> <p>You can optionally specify a named instance JDBC URL to connect to a SQL Server database. Format the string as follows:</p> <pre>jdbc:informatica:sqlserver://<database host name>\<named instance name>;databaseName=<database name></pre> <p>You can also specify a Windows authentication connection string to connect to a SQL Server database using Windows authentication. Format the string as follows:</p> <pre>jdbc:informatica:sqlserver://<database host name>:<database port>;DatabaseName=<database name>;SnapshotSerializable=true;AuthenticationMethod=ntlmjava;Domain=<SQL Server domain name></pre> <p>Note that the connection uses the NTLM authentication scheme.</p> |

2. Enter the following optional properties:

| Property | Description |
|------------------------|--|
| Secure JDBC Parameters | Secure database parameters if the Model repository database is secured with the SSL protocol. |
| Schema Name | Schema name in the Model repository database that contains monitoring data. |
| Tablespace Name | If the Model repository database is an IBM DB2 database, you can specify the name of the tablespace that contains monitoring data. |

3. Click **Test Connection** to verify the connection settings.

Configure the Collector Schedule

You can configure a custom schedule for the collector. The schedule you create overrides the default collector schedule.

Enter the following properties:

| Property | Description |
|-------------------|---|
| Repeats | The interval at which to repeat collection. |
| Repeats Frequency | The frequency at which to perform collection. The frequency is based on the repetition value you select. For example, to collect data every two hours, select Hourly as the repetition, and then set the frequency value to 2. |
| Starts on | The date and time the custom schedule takes effect. |
| Timezone | The timezone the schedule is based on. |

Configure the PowerCenter Repository Collector

Configure the PowerCenter Repository collector, which collects and uploads runtime workflow and session metrics from PowerCenter repositories within the domain. You must select PowerCenter in the list of products used by the domain on the Domain Connection panel to configure the collector.

You must configure the connection to each PowerCenter repository database in the domain.

The default collection frequency is every 1 hour. You can create a custom schedule if needed to better suit your requirements.

The collector is disabled by default. Select the **Enabled** checkbox to enable the collector.

Collecting historical data

You can configure the collector to populate Operational Insights with up to 60 days of historical data.

Historical data is collected for the previous 30 days by default. However you can specify any number of days between 1 and 60.

Data collection begins at the time the domain is added to Operational Insights. Roughly 24 hours worth of data is collected every hour, meaning that approximately 30 hours are required to populate Operational Insights with data for the prior month.

Historical data collection is enabled by default. Clear the **Collect Historical Data** checkbox to disable historical data collection.

Adding a PowerCenter repository

If the domain is a PowerCenter domain, configure a connection to each PowerCenter repository database within the domain. Supply the JDBC connection string used to connect to the PowerCenter repository database. You can optionally provide parameters required to connect to a secure database.

1. Click **Add PowerCenter Repository**.

2. Enter the following required properties:

| Property | Description |
|------------------------|--|
| Database Type | PowerCenter repository database type. |
| Service Name | Name of the PowerCenter Repository Service that manages the PowerCenter repository database. |
| Username | User name for the PowerCenter repository database. |
| Password | Password for the PowerCenter repository database. |
| JDBC Connection String | <p>JDBC connection string used to connect to the PowerCenter repository database.</p> <p>You can optionally specify a named instance JDBC URL to connect to a SQL Server database. Format the string as follows:</p> <pre>jdbc:informatica:sqlserver://<database host name>\<named instance name>;databaseName=<database name></pre> <p>You can also specify a Windows authentication connection string to connect to a SQL Server database using Windows authentication. Format the string as follows:</p> <pre>jdbc:informatica:sqlserver://<database host name>:<database port>;DatabaseName=<database name>;SnapshotSerializable=true;AuthenticationMethod=ntlmjava;Domain=<SQL Server domain name></pre> <p>Note that the connection uses the NTLM authentication scheme.</p> |

3. Enter the following optional properties:

| Property | Description |
|------------------------|--|
| Secure JDBC Parameters | Secure database parameters if the PowerCenter repository database is secured with the SSL protocol. |
| Schema Name | <p>Schema name in the PowerCenter repository database that contains monitoring data.</p> <p>To find this value in the Administrator tool, select the Services and Nodes view, and then select the PowerCenter Repository Service instance in the Domain Navigator.</p> |
| Table Name | <p>Table in the PowerCenter repository database that contains monitoring data.</p> <p>To find this value in the Administrator tool, select the Services and Nodes view, and then select the PowerCenter Repository Service instance in the Domain Navigator.</p> |

4. Select the **Enable Repository** checkbox to enable the collector to collect data from the repository.
5. Click **Test Connection** to verify the connection configuration.
6. Click **Save** to save the connection details.
7. Repeat this process for each PowerCenter repository in the domain.

Configure the Collector Schedule

You can configure a custom schedule for the collector. The schedule you create overrides the default collector schedule.

Enter the following properties:

| Property | Description |
|-------------------|---|
| Repeats | The interval at which to repeat collection. |
| Repeats Frequency | The frequency at which to perform collection. The frequency is based on the repetition value you select. For example, to collect data every two hours, select Hourly as the repetition, and then set the frequency value to 2. |
| Starts on | The date and time the custom schedule takes effect. |
| Timezone | The timezone the schedule is based on. |

Configure the Data Engineering Integration collector

Configure the Data Engineering Integration collector to collect and upload statistics on Hadoop clusters used by the domain, including statistics on jobs run on the clusters. You must select Data Engineering Integration in the list of products used by the domain on the Domain Connection panel to configure the collector.

The default collection frequency is every 1 hour. You can create a custom schedule to suit your requirements.

The collector is enabled by default. Clear the **Enabled** checkbox to disable the collector.

Click **Finish** when you finish configuring the collector.

Note: To collect statistics for a Data Engineering Integration version 10.2.2 domain, you must apply EBF-14386 to each Data Engineering Integration node in the domain.

Collecting historical data

You can configure the collector to populate Operational Insights with up to 60 days of historical data.

Historical data is collected for the previous 30 days by default. However you can specify any number of days between 1 and 60.

Data collection begins at the time the domain is added to Operational Insights. Roughly 24 hours worth of data is collected every hour, meaning that approximately 30 hours are required to populate Operational Insights with data for the prior month.

Historical data collection is enabled by default. Clear the **Collect Historical Data** checkbox to disable historical data collection.

Selecting the cluster configuration

If the domain is a Data Engineering Integration domain, select the cluster configuration used by the domain to connect to the Hadoop cluster. The Data Engineering Integration collector uses the cluster configuration to gather job execution statistics and operational metrics for the cluster.

You can view the cluster configurations created in the domain in the **Connections** tab in Informatica Administrator (the Administrator tool).

1. Click **Select Cluster Configuration**.
2. Select the cluster configuration to use to connect to the Hadoop cluster from the menu.
3. Select the **Enable Cluster Configuration** checkbox to enable the collector to collect data from the cluster.
4. To connect to a secure cluster, click **TLS Enabled**, and then specify the path and password for the cluster truststore file.
5. Click **Save** to save the configuration.

Configure the Collector Schedule

You can configure a custom schedule for the collector. The schedule you create overrides the default collector schedule.

Enter the following properties:

| Property | Description |
|-------------------|---|
| Repeats | The interval at which to repeat collection. |
| Repeats Frequency | The frequency at which to perform collection. The frequency is based on the repetition value you select. For example, to collect data every two hours, select Hourly as the repetition, and then set the frequency value to 2. |
| Starts on | The date and time the custom schedule takes effect. |
| Timezone | The timezone the schedule is based on. |

Connecting to a cluster secured using Kerberos authentication

If the Data Engineering Integration collector collects analytics from a cluster secured using Kerberos authentication, you must add custom Kerberos properties to the configuration for the Secure Agent the Data Engineering Integration domain uses.

To find the Secure Agent the collector uses,

1. Log in to Operational Insights.
2. Select the domain, and then click the **Details** tab.
3. Locate the name of the Secure Agent the domain uses in the Secure Agent Group property.
4. Click **Secure Agents** in the left hand navigation bar.
5. Select the Secure Agent, then click **Manage**.

The Details page for the Secure Agent opens in the Administrator application.

6. Click **Edit**.

- Click the **+** symbol next to a property in the Custom Configuration section of the page to add a new custom property.
- For each property, select **OpsInsights Data Collector** from the Service menu, and then select **OpsInsights** from the Type menu.
- Enter the following custom properties. The table below describes the properties to add:

| Name | Value |
|--------------------|--|
| kerberosPrincipal | The Service Principal Name (SPN) assigned in Active Directory to the user that runs Data Integration Service jobs on the cluster. |
| kerberosKeyTabFile | The path and file name of the keytab file on the node where the Secure Agent runs. On both Linux and Windows hosts, specify the value as follows: <code>/<Secure Agent installation directory>/<file name>.keytab</code> |
| kerberosConfFile | The path to the krb5.conf file on the node where the Secure Agent runs. On both Linux and Windows hosts, specify the value as follows: <code>/<Secure Agent installation directory>/krb5.conf</code> |

- Click **Save**.

Configure the Data Quality Collector

Configure the Data Quality collector to collect and upload statistics on Data Quality jobs run on clusters. You must select Data Quality in the list of products used by the domain on the Domain Connection panel to configure the collector.

The default collection frequency is every 1 hour. You can create a custom schedule to suit your requirements.

The collector is enabled by default. Clear the **Enabled** checkbox to disable the collector.

Click **Finish** when you finish configuring the collector.

Configure the Collector Schedule

You can configure a custom schedule for the collector. The schedule you create overrides the default collector schedule.

Enter the following properties:

| Property | Description |
|-------------------|---|
| Repeats | The interval at which to repeat collection. |
| Repeats Frequency | The frequency at which to perform collection. The frequency is based on the repetition value you select. For example, to collect data every two hours, select Hourly as the repetition, and then set the frequency value to 2. |

| Property | Description |
|-----------|---|
| Starts on | The date and time the custom schedule takes effect. |
| Timezone | The timezone the schedule is based on. |

Finalize the on-boarding configuration

After you complete the on-boarding process, the collectors begin collecting and uploading data to the Informatica Intelligent Cloud Services for use by Operational Insights.

Click **Save** to complete the domain on-boarding process and begin collecting data.

Search for domains

You can search for specific domains using attributes or tags assigned to domains as search parameters.

Specify domain attributes as key:value pairs. For tags, just supply the tag value. Separate multiple parameters with a comma.

For example, enter this query to search for a domain in London that is assigned the tag "Production":

```
loc:London,Production
```

The domains matching your search criteria dynamically appear in the page.

Editing or unregistering a domain

You can edit the registration details for a domain. You can also unregister a domain.

When you unregister a domain, all of the collected operational data is deleted and cannot be recovered. If auto-scaling is enabled for the domain, the auto-scaling configuration is also deleted. However, any elastic nodes running in the cloud are not removed. You must manually remove the elastic nodes from the cloud.

1. Click **All Infrastructure** in the navigation bar on the left side of the page.
2. Click the domain.
3. Click the **Monitor** tab or the **Details** tab.
4. From the edit menu, select **Edit Domain** or **Unregister Domain**.


Set the time zone

Operational Insights stores time in the Coordinated Universal Time (UTC) format.

When you log in, Operational Insights converts the time and displays it in the time zone associated with your user profile. You can edit the user profile and select a time zone based on your requirement. If you do not set the user profile time or the time is not available, the time zone is displayed in PDT.

When you hover over the data in a chart in all the Operational Insights pages, the statistics summary displays the data in the time zone set in the user profile. Additionally, tooltips also display the time range for Data Engineering Integration, Data Quality, and PowerCenter jobs.

For example, in the following image, you can view a summary of the total rows loaded for the last four weeks or the last six months for all services and domains in your infrastructure in the time zone set in the user profile:



CHAPTER 10


Monitor Data Engineering Integration domains

You can use Operational Insights to monitor your Data Engineering Integration domains. Use the **Overview** page to view job execution statistics and data processing trends for a specific Data Engineering Integration domain.

Complete the following steps to view the page:

1. Click **Data Engineering Integration** in the left hand navigation bar.
2. Click a domain on the **Domains** panel.
3. Click the **Overview** tab.

The upper part of the page displays job execution and data processing statistics for the selected date range, Data Integration Service instances, and clusters. By default, the table displays data for mapping and workflow jobs run by Data Integration Service instances submitting jobs to registered clusters over the last seven days.



The following table lists the tasks you can perform from this page:

| Task | Description |
|------|--|
| 1 | Use filters to select the date range, Data Integration Service instances, and Hadoop clusters to view data for. The data displayed in the page is updated based on your filter settings. By default, data for the last seven days is shown for all Data Integration Service instances and Hadoop clusters in the domain. |
| 2 | In the Total Instances chart, choose whether to view jobs by job completion status, job type, or execution engine. Hover over a bar segment in the chart, and then click View Jobs to view details. |
| 3 | Move your cursor across the Total Data Read and Total Rows Read charts to view details for the time range specified in the filter. |
| 4 | Click a link in the Unique Jobs region to view details for unique workflows and mappings for the selected date range, services and clusters. See " Viewing Data Engineering Integration job analytics " on page 122 for additional details. |

Viewing Data Engineering Integration job analytics

Use Data Engineering Integration job execution statistics to assess job execution performance, identify failed and long-running jobs, and troubleshoot issues.

Use filters to drill down on the job execution data you want to view. The data displayed in each page is based on the combined filters you set.

Viewing job execution data

Use summary data collected on job run executions to gain insight into job run performance. You can view the overall performance of jobs for specified time ranges, by Data Integration Service, and by job type. You can also identify jobs with increasing failure rates or jobs with increasing run times.

1. Click a domain.
2. Click the **Jobs** tab.
3. Use filters to select the jobs to view data for.

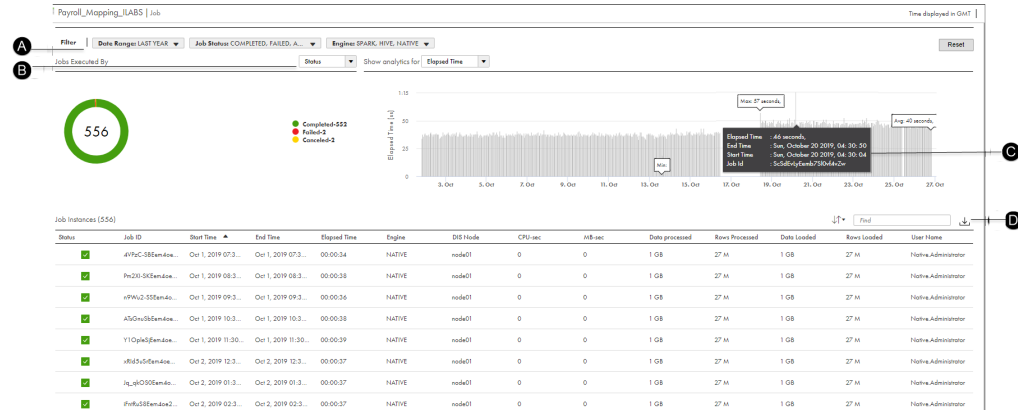
The table updates according to your filter settings. By default, the table displays data for all jobs that ran over the last seven days.

4. Click a job in the table to view details.

A page displaying statistics for individual job instances loads. Use filters to select the data that you want to view. You can select additional columns to filter on from the **Add Field** menu.

- Click a job in the table.

A graph displaying job run statistics loads in the page. You can view statistics by elapsed run time, amount of data processed, the amount of data read from source rows read, and the amount of data written to target rows. The charts display data based on the filters you set.



- Use filters to select the job data you want to view.
- Choose whether to view jobs by status or by execution engine.
- Move your cursor across the graph to view specific details. You can zoom in on the graph to view details for a specific time frame. See [“Zooming in on graph details” on page 30](#) for details.
- Click the icon to download the table data to a comma-separated value (.csv) file.

Creating a Data Engineering Integration project

You can create projects in Operational Insights to help you monitor Data Engineering Integration assets.

A project is a grouping of deployed services and applications in a Data Engineering Integration domain. You can view data processing analytics and job run statistics for the services and applications in each project you create.

- Click **Data Engineering Integration** in the left hand navigation bar.
- Click the **Projects** tab.

The Projects page displaying Data Quality projects appears. The panel for each project displays summary data for the jobs run during the last seven days.

- Click **Create Project**.
- Enter a name for the project.
The application displays the name on the Projects page.
- Add or select tags to help users search for the project.
- Select the domain to include in the project.
- Select a service within the domain to include in the project.
Expand the service to select individual applications.

CHAPTER 11


Monitor Data Quality domains

You can use Operational Insights to monitor your Data Quality domains. Use the **Overview** page to view job execution statistics and data processing trends for a specific Data Quality domain.

Complete the following steps to view the page:

1. Click **Data Quality** in the left hand navigation bar.
2. Click a domain on the **Domains** panel.
3. Click the **Overview** tab.

The upper part of the page displays job execution and data processing statistics for the selected date range and Data Integration Services. By default, the table displays data for mapping and workflow jobs run by Data Integration Services submitting jobs to clusters over the last seven days.



The following table lists the tasks you can perform from this page:

| Task | Description |
|------|--|
| 1 | Use filters to select the date range, Data Integration Service instances, and Hadoop clusters to view data for. The data displayed in the page is updated based on your filter settings. By default, data for the last seven days is shown for all Data Integration Service instances and Hadoop clusters in the domain. |
| 2 | In the Total Instances chart, choose whether to view jobs by job completion status, job type, or execution engine. Hover over a bar segment in the chart, and then click View Jobs to view details. |
| 3 | Move your cursor across the Total Data Read and Total Rows Read charts to view details for the time range specified in the filter. |
| 4 | Click a link in the Unique Jobs region to view details for unique workflows and mappings for the selected date range, services and clusters. See " Viewing Data Quality job analytics " on page 125 for additional details. |

Viewing Data Quality job analytics

Use job execution statistics to assess job execution performance, identify failed and long-running jobs, and troubleshoot issues.

Use filters to drill down on the job execution data you want to view. The data displayed in each page is based on the combined filters you set.

Viewing job execution summary data

Use summary data collected on job run executions to gain insight into job run performance. You can view the overall performance of jobs for specified time ranges, by Data Integration Service, and by job type. You can also identify jobs with increasing failure rates or jobs with increasing run times.

1. Click the **Overview** tab.

2. Click a domain.

You might need to first select a location, then select a domain within the location.

3. Click the **Jobs** tab.

A page displaying job runs loads.

4. Use filters to select the job data to view.

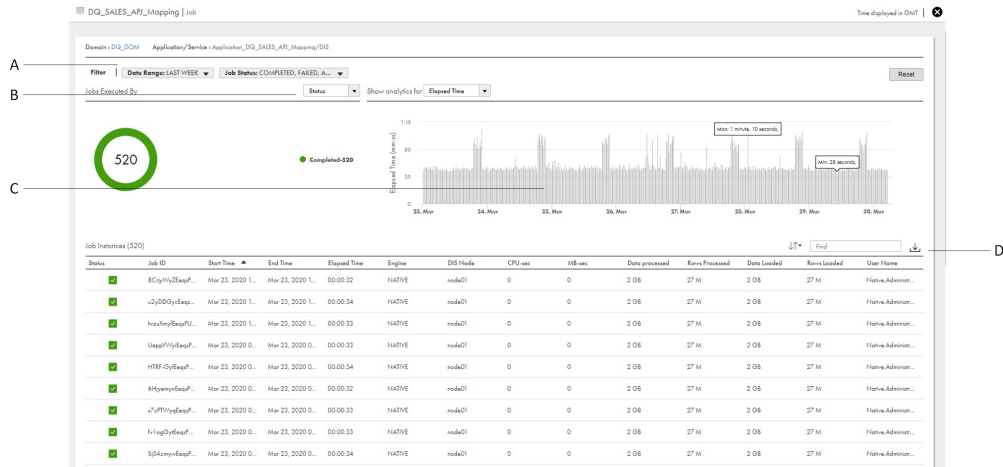
The page updates according to your filter settings. By default, the page displays data for all jobs run over the last seven days.

5. Click a job in the table to view details.

A page displaying statistics for individual job instances loads. Use filters to select the data you want to view. You can select additional columns to filter on from the **Add Field** menu.

- Click a job in the table.

A graph displaying job run statistics loads in the page. You can view statistics by elapsed run time, amount of data processed, the amount of data read from source rows read, and the amount of data written to target rows. The charts display data based on the filters you set.



- Use filters to select the job data you want to view.
- Choose whether to view jobs by status or by execution engine.
- Move your cursor across the graph to view specific details. You can zoom in on the graph to view details for a specific time frame. See [“Zooming in on graph details” on page 30](#) for details.
- Click the icon to download the table data to a comma-separated value (.csv) file.

Creating a Data Quality project

You can create projects in Operational Insights to help you monitor Data Quality assets.

A project is a grouping of deployed services and applications in a Data Quality domain. You can view data processing analytics and job run statistics for the services and applications in each project you create.

- Click **Data Quality** in the left hand navigation bar.
- Click **Projects** tab.

The Projects page displaying Data Quality projects appears. The panel for each project shows summary data for the jobs run during the last seven days.

- Click **Create Project**.

- Enter a name for the project.

The application shows the name on the Projects page.

- Add or select tags to help users search for the project.
- Select the domain to include in the project.
- Select a service within the domain to include in the project.

Expand the service to select individual applications.

CHAPTER 12

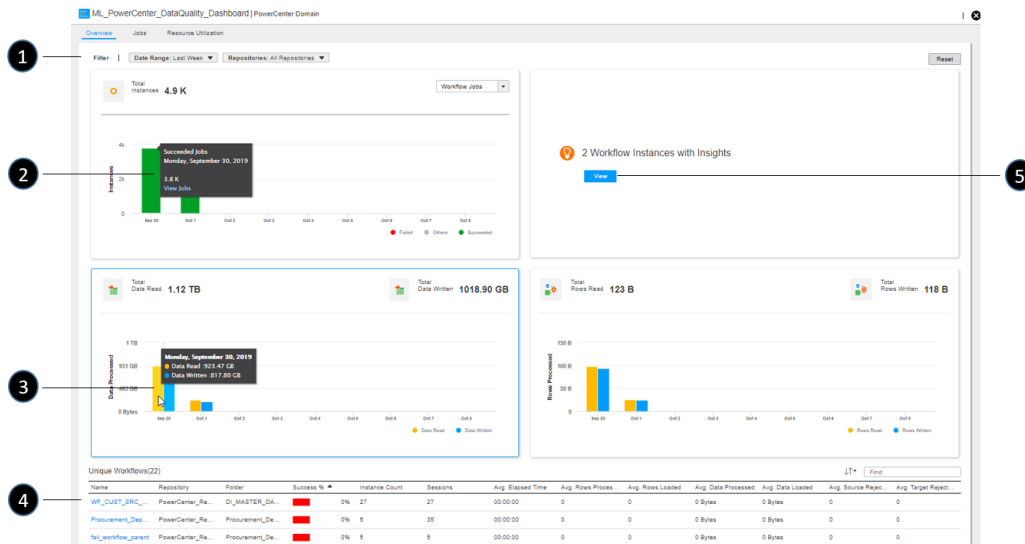
Monitor PowerCenter domains

You can use Operational Insights to monitor your PowerCenter domains. Use the **Overview** page to view workflow statistics and data processing trends for a specific domain.

Complete the following steps to view the **Overview** page:

1. Click **PowerCenter** on the Operational Insights navigation bar.
2. Click a domain on the **Domains** panel.
3. Click the **Overview** tab.

The **Overview** page displays workflow execution and data processing statistics for the selected date range and repositories and folders. You can view data based on data volume processed and on total rows moved. By default, the page displays the data volume processed for workflow jobs run over the last seven days.



The following table lists the tasks you can perform from this page:

| Task | Description |
|------|--|
| 1 | Use filters to select the date range and repositories and folders to view data for. The data displayed in the Overview page is updated based on your filter settings. |
| 3 | In the Total Instances chart, choose whether to view data based on workflow jobs or session tasks. Hover over a bar segment in the chart, and then click View Jobs to view details. |

| Task | Description |
|------|---|
| 3 | Move your cursor across the Total Data Read and Total Rows Read charts to view details for the time range specified in the filter. |
| 4 | Click a workflow in the Unique Workflows region to view detailed workflow instance statistics for the selected date range and repositories. See "Viewing PowerCenter workflow analytics" on page 128 for details. |
| 5 | Click View to display a list of workflow instances in which the CLAIRE engine has detected anomalous or abnormal behavior. See "Viewing anomalous workflow run behavior" on page 129 for details. |

Viewing PowerCenter workflow analytics

Use PowerCenter workflow statistics to assess workflow instance performance, identify failed and long-running workflow runs, and troubleshoot issues.

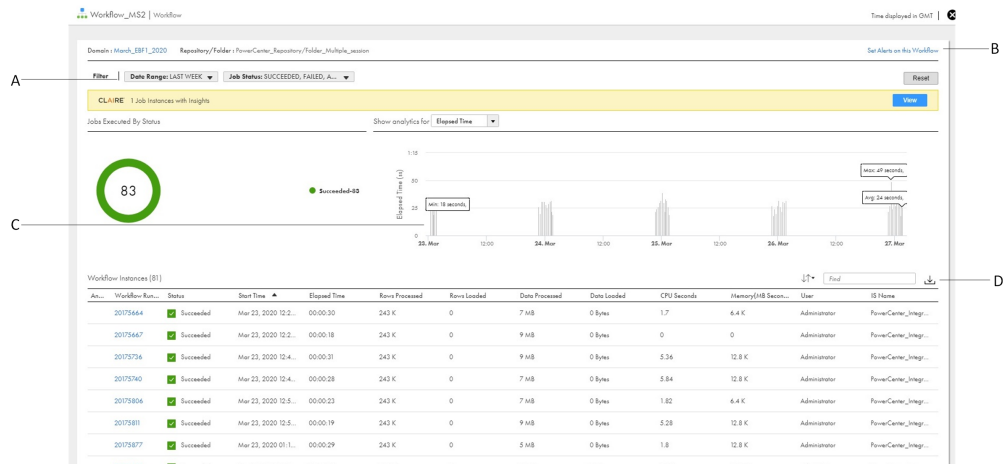
Use filters to drill down on the data you want to view. The data displayed in each page is based on the combined filters you set.

Viewing PowerCenter workflow execution data

Use data collected on workflow run executions to gain insight into job run performance. You can view the overall performance of workflows for specified time ranges, for specific repositories and folders, by instance count, and by average time elapsed and data processed. You can also identify jobs with increasing failure rates or jobs with increasing run times.

1. Click **PowerCenter** in the left hand navigation bar.
The PowerCenter dashboard page displays.
2. Click a PowerCenter domain in the dashboard page.
3. Click the **Jobs** tab.
A page displaying workflow runs displays.
4. Use filters to select the workflow data to view. You can select additional columns to filter on from the **Add Field** menu.
The table page according to your filter settings.
5. Click a workflow in the table to view details on workflow instances.

A graph displaying job run statistics loads in the page. You can view statistics by elapsed run time, amount of data processed, the amount of data read from source rows read, and the amount of data written to target rows. The charts display data based on the filters you set.



- A. Use filters to select the workflow data you want to view.
 - B. Configure alerts for the workflow.
 - C. Move your cursor across the graph to view specific details. You can zoom in on the graph to view details for a specific time frame. See [“Zooming in on graph details” on page 30](#) for details.
 - D. Click the icon to download the Workflow Instances table data to a comma-separated value (.csv) file.
- If anomalies are detected for workflow runs, you can view a list of the workflow instances that have anomalous behavior.
6. Click a Run ID in the Workflow Instances table to view task details.

Viewing anomalous workflow run behavior

Operational Insights leverages the CLAIRE engine, which employs statistical and machine learning approaches to detect data outliers and anomalies, to provide insights that notify you about abnormal PowerCenter workflow run behavior.

CLAIRE detects anomalies by analyzing elapsed run times, data processed and loaded, and rows processed and loaded for workflows each day. You can use this data to identify time periods during which anomalous or abnormal behavior occurred and determine the root cause.

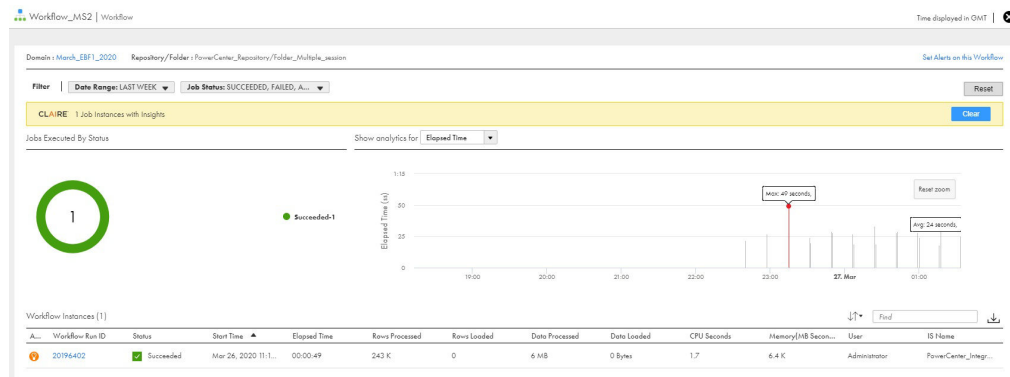
1. Click **PowerCenter** in the left hand navigation bar.
2. Perform one of the following steps:
 - Click the **New Insights** link in the panel for a domain.
 - Click a PowerCenter domain in the dashboard page, and then click **View** in the panel indicating that workflows with insights have been detected.

The Jobs page opens with a list of the workflow instances that have anomalous behavior.

3. Anomalous values are enclosed in red boxes. Click a value enclosed in a red box, and then click **Show All Anomalies of this Job**.

Click the tab above the graph for the metrics containing the anomalous run data, such as **Data Processed**.

The graph updates with data for the workflow instance. Bars in the graph correlating to anomalous behavior are indicated by a red dot.



4. Click a red dot in the graph to view the details for the anomaly.

Viewing recommendations

Use insights in the form of recommendations to improve performance, resolve errors, and avoid potential issues within PowerCenter domains.

Error remediation recommendations are generated daily for all PowerCenter domains and PowerCenter projects across the enterprise. The recommendations shown are for the most frequently occurring errors over the past 7 days. Recommendations include links to Informatica Knowledge Base articles related to the error code reported in the recommendation.

To view recommendations, complete the following steps:

1. Click **PowerCenter** in the left hand navigation bar.
2. Click the **Insights** tab.
3. Use filters to select the domain, date range and status you want to view recommendations for. You can also enter one or more comma-separated Error ID values to filter by error code.

The page displays up to 25 recommendations. Click **View More** to view the next set of recommendations.

4. Click **View** in a recommendation card to view the workflows impacted by the error. A dialog box listing the workflows opens. Click a workflow to view details.
5. Click **View More** within a recommendation card to view additional details. You can expand multiple recommendations at the same time.
6. Rate the recommendation.
 - If you click the thumbs up icon, the number of Likes is increased by 1.
 - If you click the thumbs down icon, a comment dialog opens so you can explain why you gave the recommendation a poor rating. Your feedback is not visible to other Operational Insights users. It is used by Informatica to improve the quality or usefulness of the recommendation.
7. Indicate the status of the recommendation to help you track your progress in resolving the error.

If a recommendation is not applicable to you, select the **Dismiss** status to remove the recommendation from the list.

8. Click **Read article here** to open an Informatica Knowledge Base article related to the error code in a new browser.

Viewing the resource utilization heat map

Use the heat map to quickly identify resource contention issues within a PowerCenter domain, and to analyze bottlenecks caused by too many workflow jobs running within the same time period.

The heat map provides a calendar view showing periods of heaviest and lightest resource consumption. You can view data based on memory consumption or on CPU usage. The heat map displays CPU usage data by default.

Each tile in the calendar represents a four-hour time period. The darkest colored tiles represent periods of 71% or higher consumption; the lighter colors represent periods with lower consumption. Click a tile to load details for the jobs that ran within the time period.


You can download the table containing the data for a selected tile to a comma-separated value (.CSV) file. If you note that consumption is high for a time period on one day, but lighter for the same time period on a different day, you might want to download the tables for both tiles and compare the data in each to identify possible causes.

Before you enable the resource utilization heat map for a PowerCenter domain, you need to either apply the applicable EBFs or set the custom properties. For more information, see the following Knowledge Base article: [563791](#)

To view the resource utilization heat map, complete the following steps:

1. Click **PowerCenter** in the left hand navigation bar.
The PowerCenter dashboard page displays.
2. Click a PowerCenter domain in the dashboard page.
3. Click the **Overview** tab.
4. Click a PowerCenter domain.

- Click the **Resource Utilization** tab.



- Choose to display CPU usage data or memory consumption data.
- Click the tile for a time period. Details for the jobs that ran within the time period load in the table below.
- Click a Run Id to view the job run task details.
- Click to download table data to a comma-separated value (.csv) file.

Viewing domain resource utilization


Use the Domain Resource Utilization graph to view resource consumption statistics for a domain. You can zoom in on a graph to view details for a specific time frame.

- Click **All Infrastructure**.
- Click the **Monitor** tab.
- Click a domain.

The Domain Resource Utilization graph appears in the lower portion of the page.

- Select the time period to view details for in the graph. You can choose to view statistics for the last 24 hours, the last week, or the last month.
- Move your cursor across the graph to view specific details.
- To zoom in on graph details, place your cursor on the time frame start point in a graph.

7. Left-click at the start point and drag your cursor to the end point, as shown in the following image:



The resource utilization graphs update to display data only for the specified time frame.

8. Click **Reset zoom** to return the graph to the original state.

Creating a PowerCenter project

You can create projects in Operational Insights to help you monitor PowerCenter assets.

A project is a grouping of repositories and folders in a PowerCenter domain. You can view data processing analytics and workflow run statistics for the repositories and folders for each project you create.

Operational Insights leverages the CLAIRE engine, which employs statistical and machine learning approaches to detect data outliers and anomalies, to detect abnormal PowerCenter workflow run behavior within a project. You can configure the application to send email notifications when anomalous behavior is detected. You can specify the users and groups of users that receive anomaly alerts, or allow all project users to receive alerts by default.

1. Click **PowerCenter** in the left hand navigation bar.
2. Click **Projects** tab.

The Projects page displaying all PowerCenter projects appears. The panel for each project displays summary data for the workflows run during the last seven days.

3. Click **Create Project**.
4. Enter a name for the project.
The application displays the name on the Projects page.
5. Add or select tags to help users search for the project.
6. Select the domain to include in the project.
7. Select the PowerCenter repository folders within the domain to include in the project.


Expand a repository to select individual folders.

8. Optionally enable alerts, and then specify the users or groups of users that receive email notifications.

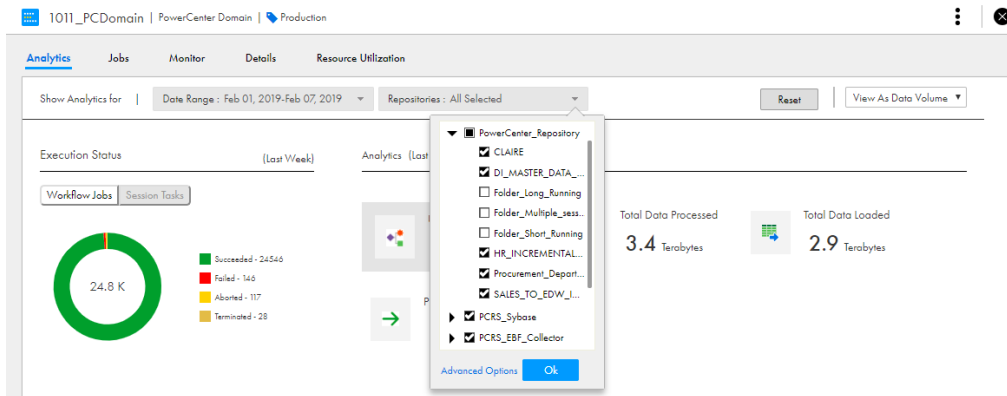
Using PowerCenter repository filters

Use PowerCenter repository filters to display run-time workflow analytics collected during specific time periods. You can also select specific PowerCenter repositories and folders within each repository to display data for. You can use advanced filtering to further refine results.

1. Select the date range you want to show statistics for from the Date Range menu.



2. Select the repositories you want to display data for from the Repositories menu. All repositories are selected by default; clear those you don't want to include.



3. Optionally select the folders within each repository you want to view data for. All folders are selected by default; clear those you don't want to include.
You can also click **Advanced Options** to select the folders within each repository you want to view data for.
4. Click **OK**.
5. Click **Apply** to reload the page with analytics data for the selected date range and repositories.

PowerCenter alerts

You can configure Operational Insights to send email notifications to Informatica Intelligent Cloud Services users, user groups, or email recipients outside the organization when an issue occurs in a PowerCenter

workflow. For example, you can set conditions to alert a user when the workflow stops running, or the CPU utilization matches the value specified in the alert condition.

On the **PowerCenter Alerts** tab, you can view the configured alerts. You can enable or disable alerts that Operational Insights monitors.

You can set the following PowerCenter alerts:

CLAIRE alerts

Alerts for anomalous or abnormal behavior in a PowerCenter project detected by the CLAIRE™ engine. You can enable CLAIRE alerts when you create a PowerCenter project, and then specify the users or groups of users that need to receive email notifications.

For more information about configuring CLAIRE alerts, see [“Creating a PowerCenter project” on page 133](#).

Workflow alerts

Alerts for issues that occur in a workflow within the PowerCenter domain or in a project created in Operational Insights.

To create an alert rule for a PowerCenter workflow, specify the alert condition for which you want to generate an alert. You can specify the repository and folder filters in the domain or project to set alerts for specific workflows. If you do not specify the repository and folder filters, the alert applies to all workflows in the domain or project.

When you set the alert, you must also specify the Informatica Intelligent Cloud Services users, user groups, or email recipients outside the organization that must receive the alert notifications when the issue matches the specified alert condition.

You can also create an alert script that Operational Insights runs to perform additional actions when an alert is triggered. For more information about creating and using alert scripts, see [“Use alert scripts” on page 33](#).

INDEX

A

alerts

- application ingestion and replication jobs [97](#)
- Data Integration jobs [48](#)
- database ingestion and replication jobs [97](#)

C

Cloud Application Integration community

URL [13](#)

Cloud Developer community

URL [13](#)

configuring alerts

- Data Integration [49](#)
- domain [32](#)
- infrastructure [32](#)
- Secure Agent [32](#)

connection properties [46](#)

D

Data Ingestion and Replication jobs

alerts

- configuring for application or database ingestion and replication jobs [98](#)

configuring alerts [98](#)

Data Integration community

URL [13](#)

Data Integration job details [43](#)

Data Integration jobs [39](#)

Data Profiling job details [70](#)

Data Profiling jobs [69](#)

directories

- configuring Secure Agent login to access [21](#)

I

Informatica Global Customer Support

contact information [14](#)

Informatica Intelligent Cloud Services

web site [13](#)

J

job details [43](#), [70](#)

Jobs data

- exporting [42](#)

L

Linux

- configuring proxy settings [25](#)
- uninstalling the Secure Agent [26](#)

M

maintenance outages [14](#)

monitoring Data Integration jobs [39](#)

monitoring Data Profiling jobs [69](#)

monitoring ingestion and replication jobs

database ingestion and replication job details [85](#)

Job Overview tab [85](#)

job properties in job lists [77](#)

monitoring all jobs from Operational Insights and Monitor [74](#)

monitoring My Jobs [73](#)

Object Details tab [85](#)

viewing job details [79](#)

monitoring jobs [39](#), [69](#)

P

proxy settings

- configuring on Linux [25](#)
- configuring on Windows [20](#)

R

requirements

- Secure Agent [17](#), [22](#)

S

scheduled jobs

Data Integration [48](#)

Secure Agent Manager

launching [17](#)

Secure Agents

configuring a Windows service login [21](#)

installing on Linux [23](#)

installing on Windows [18](#)

permissions on Linux [23](#)

permissions on Windows [17](#)

registering on Linux [23](#)

registering on Windows [18](#)

requirements on Linux [22](#)

requirements on Windows [17](#)

starting on Windows [17](#)

uninstalling on Linux [26](#)

uninstalling on Windows [21](#)

status

Informatica Intelligent Cloud Services [14](#)

streaming ingestion and replication jobs [95](#)

system status [14](#)

T

trust site

description [14](#)

U

upgrade notifications [14](#)

V

viewing connections [46](#)

viewing job details [43](#), [70](#)

W

web site [13](#)

Windows

configuring proxy settings [20](#)

Windows service

configuring Secure Agent login [21](#)