



Informatica® Data Ingestion and Replication
July 2024

Monitoring Data Ingestion and Replication Jobs

© Copyright Informatica LLC 2019, 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: 2024-07-12

Table of Contents

- Preface 4**
 - Informatica Resources. 4
 - Informatica Documentation. 4
 - Informatica Intelligent Cloud Services web site. 4
 - Informatica Intelligent Cloud Services Communities. 4
 - Informatica Intelligent Cloud Services Marketplace. 4
 - Data Integration connector documentation. 5
 - Informatica Knowledge Base. 5
 - Informatica Intelligent Cloud Services Trust Center. 5
 - Informatica Global Customer Support. 5
- Chapter 1: Monitoring Data Ingestion and Replication jobs..... 6**
 - Monitoring your ingestion and replication jobs. 7
 - Monitoring all ingestion and replication jobs. 8
 - Job properties. 11
 - Viewing details for an ingestion and replication job. 13
 - Application ingestion and replication job details. 13
 - Database ingestion and replication job details. 19
 - File ingestion and replication job details. 26
 - Streaming ingestion and replication job details. 28
- Index..... 32**

Preface

Read *Monitoring Data Ingestion and Replication Jobs* to find out how you can view and monitor ingestion and replication jobs for the tasks that you deployed or for all tasks that any members of your organization deployed.

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>.

CHAPTER 1

Monitoring Data Ingestion and Replication jobs

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 11](#).

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:

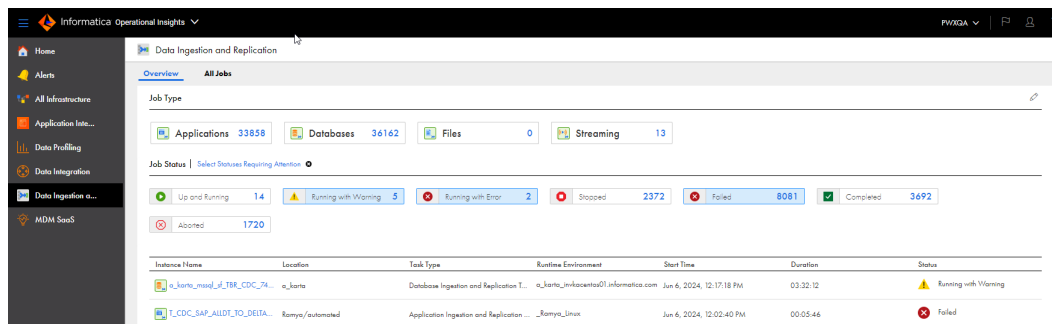


You can use the buttons at the top to filter the jobs by job type and status. If you want to see jobs of a particular job type and status, first select the job type and then select the status. You cannot select multiple status buttons at the same time. To clear a filter, click the selected button again.

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:

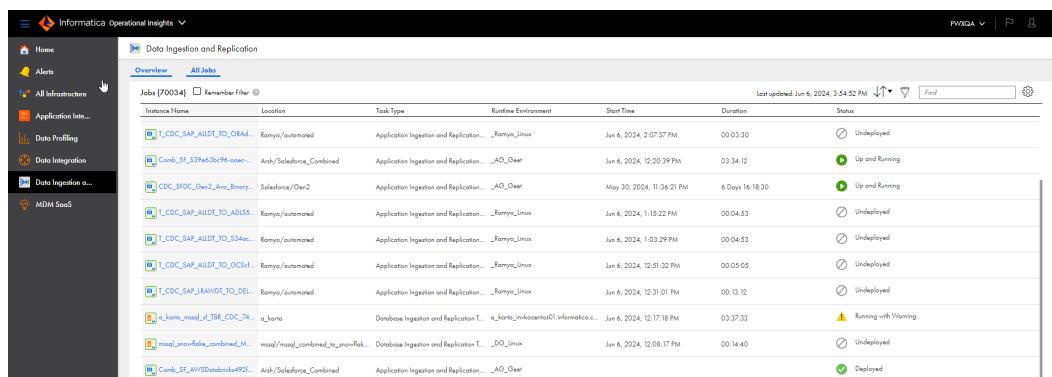


This option lists the application ingestion and replication jobs or database ingestion and replication jobs with the **Failed** or **Running with Warning** status, file ingestion and replication jobs with the **Failed** status, and streaming ingestion and replication jobs with the **Running with Error** 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.

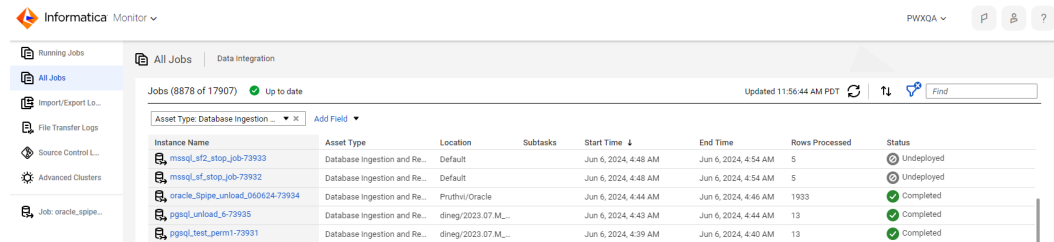
You can also use the Actions menu at the end of each row to perform an action on the job, such as Redeploy, depending on the job type and status.

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
mssql_sf2_stop_job-73933	Database Ingestion and Re...	Default		Jun 6, 2024, 4:48 AM	Jun 6, 2024, 4:54 AM	5	Undeployed
mssql_sf2_stop_job-73932	Database Ingestion and Re...	Default		Jun 6, 2024, 4:48 AM	Jun 6, 2024, 4:54 AM	5	Undeployed
oracle_splice_unload_000624-73934	Database Ingestion and Re...	Pruthvi/Oracle		Jun 6, 2024, 4:44 AM	Jun 6, 2024, 4:46 AM	1933	Completed
pgegal_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
pgegal_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: <taskname>-<job_instance_number></p> <p>You can click the instance name to view detailed information about the job.</p> <p>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 ReplicationTask- 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 13, "Database ingestion and replication job details" on page 19, "File ingestion and replication job details" on page 26 ("Results" section), or "Streaming ingestion and replication job details" on page 28.</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: <code><taskname>-<job_instance_number></code></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 (-): <code><taskname>_<job_instance_number></code>
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>
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 13, "Database ingestion and replication job details" on page 19, "File ingestion and replication job details" on page 26 ("Results" section), or "Streaming ingestion and replication job details" on page 28.</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	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.
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	The type of task, which is Application Ingestion and ReplicationTask .
Task Location	The project or project folder that contains the application ingestion and replication task definition.
Started By	The name of the user who started the job.
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	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.
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 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.</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 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.

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 or detects a DDL change.

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.

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:

The screenshot displays the job details for 'ora_gen2_combined_73735'. The interface is divided into several panes:

- Task Summary:** Shows the job status as 'Running with Warning'. It includes a diagram of the data flow from 'Source' (Oracle Database) to 'Target' (Azure Data Lake Storage).
- Overview:** Displays a donut chart showing the progress of the job. The chart indicates 12 total items, with 8.33% (1 item) in error and 91.67% (11 items) running. The duration is 1 Day 06:00:28, and the start time is Jun 5, 2024, 5:50:27 AM.
- Object Detail (12):** A table listing the objects being processed. The table has columns for Object, Target Object, Agent Name, Stage, Status, and Log. Two objects are listed: 'SREE_SRC_ALLCHAR_BK' and 'SREE_SRC_ALDTE_BK', both in 'Normal' stage and 'Running' status.
- Alerts (13):** A list of alerts generated during the job. The alerts are categorized by level (Critical, Info) and include a code and timestamp. The first alert is 'DBMIP_27008' at 'Jun 6, 2024, 4:41:30 AM'.

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	<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, Database 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 Task Summary and the Object Detail panes.</p>
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	<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	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	<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	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 Detail (0)							
Object	Target Object	Status	Log				
DBMICDC_pg013_rc_rongetype1_int	KARHEEK_pg013_rc_rongetype1_int_unload	Stopped	Select log	Inserts	Updates	Deletes	LOBs
				11	4	11	0
DBMICDC_pg014_rc_rongetype2_int	KARHEEK_pg014_rc_rongetype2_int_unload	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	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>
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.

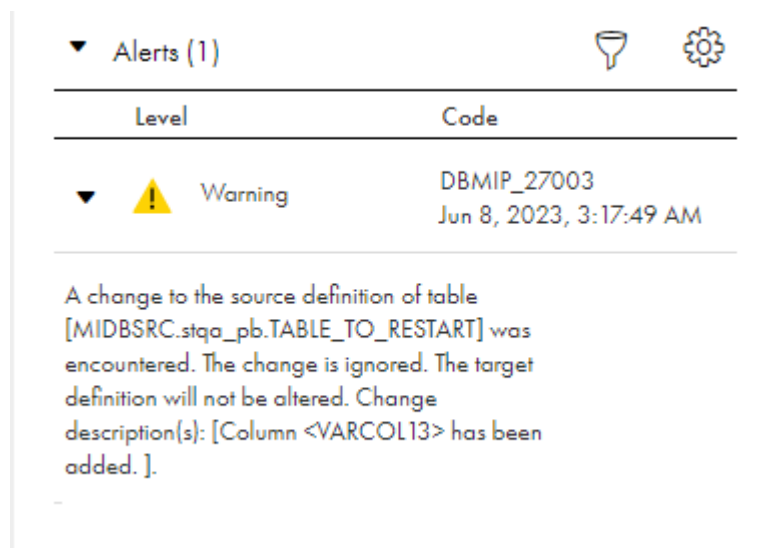
Column	Description
CDV Last Run	<p>If data validation is available for the task, this column shows the date and time when the data validation was last performed for the given subtask. If the data validation has not been run, the field value is N/A.</p> <p>This column appears in initial load jobs that 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:



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 or detects a DDL change.

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.

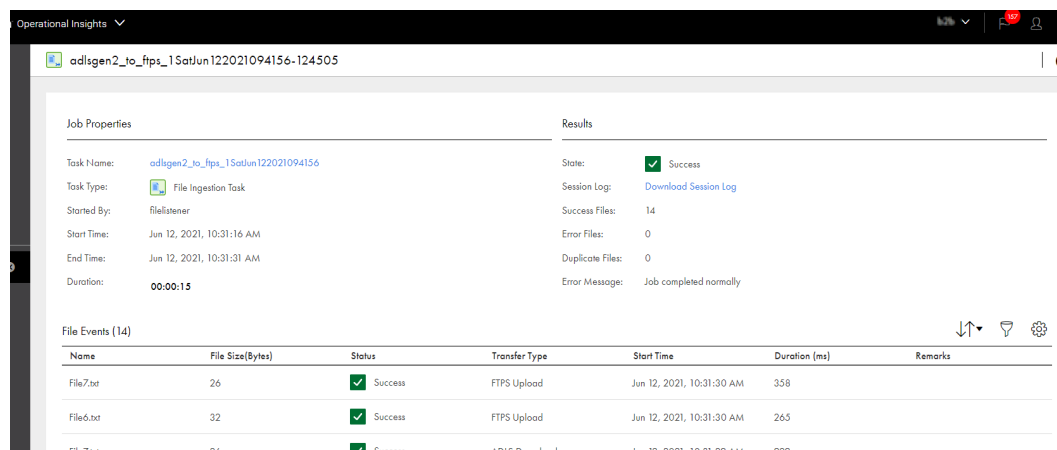
Click the expander arrow to display a description of the event.

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:



Job Properties		Results	
Task Name:	adlsgen2_to_ftps_15atjun122021094156	State:	Success
Task Type:	File Ingestion Task	Session Log:	Download Session Log
Started By:	filelistener	Success Files:	14
Start Time:	Jun 12, 2021, 10:31:16 AM	Error Files:	0
End Time:	Jun 12, 2021, 10:31:31 AM	Duplicate Files:	0
Duration:	00:00:15	Error Messages:	Job completed normally

File Events (14)						
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:</p> <pre><Secure Agent installation directory>/apps/MassIngestionRuntime/data/taskLogs</pre>
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.
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 hh:mm:ss format.
Start Time	The date and time when the job was deployed.

Property	Description
Runtime Environment	Name of the runtime environment that the job uses to run.
Download Log	<p>You can download one of the following types of logs:</p> <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. <p>To download a log to your local system, click the Download icon.</p>

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:

Overview	Alert	Performance	Past Run
Find			
Alert	Time		
⚠ KafkaToFlatfileTargetWithFormatConverterWithXML_FormatConverter : ConvertRecord[id=6ee4a30e-9ed9-4e00-87bf-d6cc0e0... Show More	Nov 24, 2021, 01:30 AM		
⚠ KafkaToFlatfileTargetWithFormatConverterWithXML_FormatConverter : ConvertRecord[id=6ee4a30e-9ed9-4e00-87bf-d6cc0e0... Show More	Nov 24, 2021, 01:04 AM		
⚠ KafkaToFlatfileTargetWithFormatConverterWithXML_FormatConverter : ConvertRecord[id=6ee4a30e-9ed9-4e00-87bf-d6cc0e0... Show More	Nov 24, 2021, 12:55 AM		
⚠ KafkaToFlatfileTargetWithFormatConverterWithXML_FormatConverter : ConvertRecord[id=6ee4a30e-9ed9-4e00-87bf-d6cc0e0... Show More	Nov 24, 2021, 12:33 AM		


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.

INDEX

C

Cloud Application Integration community
URL [4](#)
Cloud Developer community
URL [4](#)

D

Data Integration community
URL [4](#)

I

Informatica Global Customer Support
contact information [5](#)
Informatica Intelligent Cloud Services
web site [4](#)

M

maintenance outages [5](#)
monitoring ingestion and replication jobs
database ingestion and replication job details [19](#)
Job Overview tab [19](#)
job properties in job lists [11](#)
monitoring all jobs from Operational Insights and Monitor [8](#)

monitoring ingestion and replication jobs (*continued*)
monitoring My Jobs [7](#)
Object Details tab [19](#)
overview [6](#)
viewing job details [13](#)

S

status
Informatica Intelligent Cloud Services [5](#)
streaming ingestion and replication jobs [28](#)
system status [5](#)

T

trust site
description [5](#)

U

upgrade notifications [5](#)

W

web site [4](#)