



Informatica® Mass Ingestion
January 2023

Monitoring Ingestion Jobs

© Copyright Informatica LLC 2019, 2023

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

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

Informatica, 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: 2023-01-09

Table of Contents

- Preface 4**
- Chapter 1: Monitoring Mass Ingestion Jobs..... 5**
 - Monitoring your ingestion jobs. 5
 - Monitoring all ingestion jobs. 6
 - Job properties. 9
 - Viewing details for an ingestion job. 9
 - Application ingestion job details. 10
 - Database ingestion job details. 15
 - File ingestion job details. 21
 - Streaming ingestion job details. 23
- Index..... 28**

Preface

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

CHAPTER 1

Monitoring Mass Ingestion Jobs

You can monitor the progress, performance, and status of ingestion jobs from the Mass Ingestion and Operational Insights services.

Depending on the service you use and type of ingestion job, you can view following monitoring information:

- On the **My Jobs** page in the Mass Ingestion service, monitor the ingestion jobs for the ingestion tasks that you deployed. You can view a list of your jobs that includes general job properties such as the task type, runtime environment, start time, duration, and current job state.
- On the **Mass Ingestion** page in Operational Insights service, monitor *all* types of ingestion jobs that any member of your organization deployed. You can view the following types of information:
 - Summary counts of ingestion jobs by task type and job status.
 - Recent jobs that require your attention because they failed or are running with errors or warnings.
 - A list of all ingestion jobs by type, including the general job properties.
- From either the list of your jobs or the list of all jobs, you can drill down to details for a specific job by clicking the job name. You can view overview job information, source object processing details, and alerts.

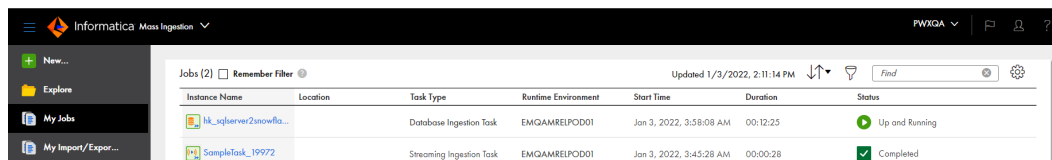
Note: Usually, a job name corresponds to the ingestion task name. For application ingestion and database ingestion jobs, the job name has the format *taskname-job_instance_number*. The number is incremented each time the job is deployed.

Monitoring your ingestion jobs

On the **My Jobs** page in Mass Ingestion, you can monitor the ingestion jobs for the tasks that you deployed. If you defined multiple types of ingestion tasks, such as application ingestion tasks, database ingestion tasks, file ingestion tasks, and streaming ingestion tasks, the list includes the jobs for all of the task types.

The **My Jobs** page shows information about each job instance, including its status.

For example, the following image shows the **My Jobs** page with a database ingestion job and a streaming ingestion job:



The screenshot shows the Informatica Mass Ingestion interface. On the left is a navigation menu with options: New..., Explore, My Jobs (selected), and My Import/Export... The main area displays a table titled 'Jobs (2)' with a 'Remember Filter' button. The table has columns: Instance Name, Location, Task Type, Runtime Environment, Start Time, Duration, and Status. It lists two jobs: 'db_sqlserver2unowfla...' (Database Ingestion Task, Up and Running) and 'SampleTask_19972' (Streaming Ingestion Task, Completed). The table is updated as of 1/3/2022, 2:11:14 PM.

Instance Name	Location	Task Type	Runtime Environment	Start Time	Duration	Status
db_sqlserver2unowfla...		Database Ingestion Task	EMGAMREIPOD01	Jan 3, 2022, 3:58:08 AM	00:12:25	Up and Running
SampleTask_19972		Streaming Ingestion Task	EMGAMREIPOD01	Jan 3, 2022, 3:45:28 AM	00:00:28	Completed

For descriptions of the columns, see [“Job properties” on page 9](#). These columns are the same as those shown for all ingestion jobs on the **All Jobs** tab of the **Mass Ingestion** page in Operational Insights.

To find a job in a long list, use any 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. The default sort order for application ingestion jobs, database ingestion jobs, and streaming ingestion jobs is the time of task deployment, from latest to earliest. The default sort order for file ingestion jobs is the job start time, from latest to earliest.
- 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. In Mass Ingestion, 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 of the listed fields. For the **Instance Name** field, you can enter the full instance name or part of the name. In Mass Ingestion, 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. In Mass Ingestion, you can save the filter for subsequent sessions by selecting the **Remember Filter** check box. To clear existing filter criteria, click the Filter icon again.

From the action (...) 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.

Monitoring all ingestion jobs

On the **Mass Ingestion** page in Operational Insights service, you can monitor all ingestion jobs that were deployed from the Mass Ingestion service, including application ingestion jobs, database ingestion jobs, file ingestion jobs, and streaming ingestion jobs.

The **Mass Ingestion** page has the following tabs:

- The **Overview** tab displays buttons that you can use to filter the list of ingestion jobs by job type and state.
- The **All Jobs** tab lists all types of ingestion jobs that any member in your organization created and deployed. It includes the same column properties as on the **My Jobs** page in the Mass Ingestion service.

Overview tab

The **Overview** tab Initially lists all types of ingestion jobs with all statuses. Each button shows the number of jobs with that job type or status. For example:

Instance Name	Location	Task Type	Runtime Environment	Start Time	Duration	Status
ms2d1_7T_m_31401	Shreddha/Pel_rntelLoad	Database Ingestion Task	mlkv7qpd01.informatix.com	May 2, 2022, 8:23:41 AM	00:02:56	Undeployed
DBMI-8652_7399	Krishna/Cricka_Target	Database Ingestion Task	ENVGAMREIPO001	May 2, 2022, 8:21:59 AM	00:28:50	Undeployed
ms2d1_7T_m_31404	Shreddha/Pel_rntelLoad	Database Ingestion Task	mlkv7qpd01.informatix.com	May 2, 2022, 8:21:22 AM	00:02:12	Completed

You can use the buttons at the top to filter the jobs by job type and status or click **Select Statuses Requiring Attention** to show only the jobs that have a status of concern. The following example shows the **Overview** tab for all jobs with statuses requiring attention:

Instance Name	Location	Task Type	Runtime Environment	Start Time	Duration	Status
CDC_HubNet_Obs_20993	Krakra/Oracle_Target	Database Ingestion Task	EMGAARBP001	Apr 27, 2022, 3:58:34 AM	00:03:24	Failed
IBC_ORACLE_ROOT_KAFKA_30588	Alia/Kafka	Database Ingestion Task	EMGAARBP001	Apr 27, 2022, 3:46:59 AM	00:02:36	Failed

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** icon. Then in the **Reorder Job Status** dialog box, select and drag a job status row up or down.

To filter the list of jobs on the **Overview** tab, use in any of the following methods:

- To see only application ingestion jobs, database ingestion jobs, file ingestion jobs, or streaming ingestion jobs, click the button for a job type. The selected button is highlighted in blue. To see all types of ingestion jobs again, click the selected button again. If you selected status filters for a job type, those filters are also cleared when you deselect the job type.
- To see the jobs that match a particular job status, click a status button. 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 on the **Overview** tab. To clear the filter, click the selected status button again.
- To see all jobs with a status that might require your attention, click **Select Statuses Requiring Attention** above the job status buttons. This option lists the application ingestion jobs or database ingestion jobs with the **Failed** or **Running with Warning** status, file ingestion jobs with the **Failed** status, or streaming ingestion 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 ingestion jobs with optional filtering by job instance name, task type, runtime environment, and status. You can also perform some actions on a job from this tab, such as Undeploy or Redeploy.

For example, the following image shows the **All Jobs** tab filtered to show application ingestion jobs:

Instance Name	Location	Task Type	Runtime Environment	Start Time	Duration	Status
NetSuite_Comb...	Krishna	Application Ingestion Task	EMQAMREIPOD01	Jan 3, 2022, 4:18:27 AM	00:04:04	Up and Running
servicesnow_in_1...	spoorthi/servicesnow	Application Ingestion Task	EMQAMREIPOD01	Jan 3, 2022, 4:08:31 AM	00:00:26	Completed
MSD_GBQ_com...	padmini/MSD365	Application Ingestion Task	imv79pom28	Jan 3, 2022, 3:58:08 AM	00:12:25	Failed
CDC_SFDC_Oen...	UpgradeTesting_Jan2022...	Application Ingestion Task	Geet_Linux	Jan 3, 2022, 3:53:10 AM	00:01:49	Failed
servicesnow_in_1...	spoorthi/servicesnow	Application Ingestion Task	EMQAMREIPOD01	Jan 3, 2022, 3:45:28 AM	00:00:28	Completed

If the list of jobs is long, use any of the following methods to make finding the job easier:

- To **sort** the listed jobs, click a column heading or click the Sort up/down arrows icon and select a field to sort by. The default sort order for application ingestion jobs, database ingestion jobs, and streaming ingestion jobs is the time of task deployment, from latest to earliest. The default sort order for file ingestion jobs is the job start time, from latest to earliest.
- 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. In Operational Insights, 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. In Operational Insights, 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. In Operational Insights, 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.

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

You can perform some actions on a job, depending on the job status and task type. From the actions (...) menu at the right end of each job row, select an action type.

Job properties

The lists of ingestion jobs on the **My Jobs** page in the Mass Ingestion service and on the **All Jobs** tab of the **Mass Ingestion** page in Operational Insights display properties for each job. The properties provide a high-level view of the job status.

The following table describes the job properties:

Property	Description
Instance Name	The generated name of the job instance in the following format: <task_name>_<instance_number> You can click the instance name to view detailed information about the job. Note: If you edit the name of the associated ingestion task, the job name remains the same.
Location	The project or project\subfolder, where the task definition associated with the job exists. For example: Myproject\Oracle Note: This property is blank for any jobs deployed prior to the Fall 2020 releases. If you move a task definition to another folder, the Location value is not updated.
Task Type	The type of ingestion task. This value must be Application Ingestion Task , Database Ingestion Task , File Ingestion Task , or Streaming Ingestion Task .
Runtime Environment	The name of the runtime environment in which the job runs.
Start Time	For application ingestion and database ingestion jobs, the date and time when the job was deployed. For file ingestion jobs, the date and time when the job started. For streaming ingestion jobs, the date and time when the job was deployed.
Duration	For application ingestion and database ingestion jobs, the amount of time that the job has run since it was deployed. For jobs that are in a Completed, Stopped, Failed, or Aborted state, the amount of time between the date and time the job was deployed and when it acquired its current state. For file ingestion jobs, the amount of time that the job has run. For streaming ingestion jobs, the amount of time that the job has been running.
Status	The current status of the job, such as Deploying, Up and Running, or Undeployed. The set of valid statuses vary by type of ingestion task. For more information, see the "Job Overview tab" section in "Database ingestion job details" on page 15 or "Streaming ingestion job details" on page 23 or the "Results" section in "File ingestion job details" on page 21 .

Viewing details for an ingestion job

On the **My Jobs** page in the Mass Ingestion service or on the **All Jobs** tab of the **Mass Ingestion** page in Operational Insights service, you can drill down on a specific ingestion job to display job details.

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 job details

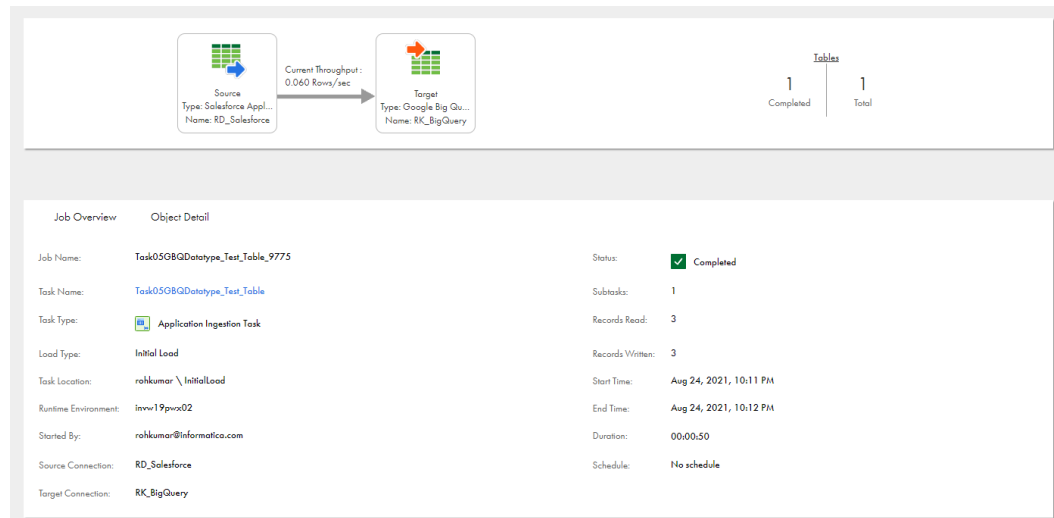
For application ingestion jobs, you can view job-specific details on the **Job Overview**, **Object Detail**, and **Alerts** tabs. To access these tabs, drill down on a job from the **My Jobs** page in Mass Ingestion or from the **All Jobs** tab on the Mass Ingestion page in Operational Insights.

Job Overview tab

On the **Job Overview** tab, view detailed information for the entire job, including the name of the associated task, the load type, source and target connection names, current status, number of records read and written, start and end times, and run duration. For incremental load jobs and combined initial and incremental load jobs, you can also download the job log.

The following image shows the **Job Overview** tab for a completed application ingestion job:

Note: The diagram at the top of the page displays the calculated data throughput, in rows per second, if the job successfully propagated 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.



The following table describes the job overview properties:

Property	Description
Job Name	The name of the job. Application ingestion job names have format <i>task name-job instance number</i> .
Task Name	The name of the associated ingestion task. You can click the task-name link to view or edit task details in Mass Ingestion, if necessary. If you edit the task, you must redeploy it for the updated task definition to be used for a job.
Task Type	The type of task, which is Application Ingestion Task .
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.

Property	Description
Task Location	The project or project folder that contains the ingestion task definition.
Runtime Environment	The name of the runtime environment that the job uses to run.
Started By	The name of the user who started the job.
Source Connection	The name of the source connection.
Target Connection	The name of the target connection.
Status	<p>The status of the job, which can be one of the following values:</p> <ul style="list-style-type: none"> - 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 Mass Ingestion Databases (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 object-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.
Subtasks	The number of subtasks that the application ingestion job used to propagate data from source objects to the target. When a job runs, it uses a separate subtask to process each source object.
Records Read	<p>The number of records that were read from the source.</p> <p>Note: The first time you run a job associated with an application ingestion combined initial and incremental load task, the Records Read count might be greater than the total number of object-level DML change records read. This behavior occurs because the initial load portion of combined processing always starts after change data capture begins. As a result, some change records are included in the Records Read count before initial load processing starts and adds more records to the count.</p>
Records Written	<p>For a Microsoft Azure Synapse Analytics target, the number of records written to the intermediate Microsoft Azure Data Lake Storage files.</p> <p>For a Snowflake target, the number of records written to the internal staging area that is created when the job runs.</p>
Start Time	The date and time when the job was deployed.
End Time	The date and time when the job ended because it completed processing, was stopped, or failed. This field is not displayed for running jobs
Duration	The amount of time, in the hh:mm:ss format, that the job ran before it ended.

Property	Description
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.
Log	<p>For incremental load jobs and combined initial and incremental load jobs, you can download the job execution log for the entire job run. Select one of the following log types:</p> <ul style="list-style-type: none"> - 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. <p>To download a log to your local system, click the Download icon.</p> <p>Note: For initial load jobs, you can get the job log for a specific source object from the Object Detail tab.</p>

Object Detail tab

On the **Object Detail** tab, view statistics and status information by source object from the last run of an application ingestion job.

The following image shows the **Object Detail** tab for an application ingestion job:

Source: Type: Salesforce Appl... Name: RD_Salesforce

Current Throughput: 0.060 Rows/sec

Target: Type: Google Big Qu... Name: RK_BigQuery

Tables: 1 Completed, 1 Total

Job Overview | Object Detail

Items (1) Updated 9/30/2021, 3:00:53 AM

Table Name	Target Object	Records Read	Records Written	Task Duration	Status	Log
SALESFORCE.Datatype_Test_Table...	ELENA_TEST.Datatype_Test_Table...	3	3	00:00:33	Completed	Select log [Download Icon]

The following table describes the properties that are displayed for each object:

Column	Description
Table Name	<p>The name of the source object 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 LOBs, Inserts, Deletes, Updates, 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 table that is mapped to the source object.
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>This information is available only at the job-level on the Job Overview tab.</p> <p>For a Microsoft Azure Synapse Analytics target, the number of records written to the intermediate Microsoft Azure Data Lake Storage files.</p> <p>For a Snowflake target, the number of records written to the internal staging area that is created when the job runs.</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 object is in the Error on Retry state, the next attempt to process the object has not yet started. - Started. Initial load processing has started. - Unloading. The subtask is unloading data from the object as part of initial load processing. - Unloaded. The subtask has finished unloading data from the object as part of initial load processing. - Completed. The subtask completed initial load processing of the object. - Normal. The subtask completed initial load processing of the object and has started CDC processing of the object. - Cancelled. Initial load processing was cancelled or stopped. - Error. The subtask detected an error in the source table.

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 Mass Ingestion Databases (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>For initial load jobs, you can download the job execution log for a source object. Select one of the following log types:</p> <ul style="list-style-type: none"> - Complete Log. The complete log for the object subtask from job execution. - Error. The log that contains error messages. This log type is available only for a Failed subtask. <p>For incremental load jobs, you can get the complete log for the entire job run from the Job Overview tab.</p> <p>For combined initial and incremental load jobs, you can download the Stage Log. This log covers the transition from initial to incremental loading for a source object.</p> <p>To download a log locally, 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>
Actions menu > Resync	<p>For a subtask in a combined initial and incremental load job, if the subtask stage is Normal and the subtask state is any state 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 Mass Ingestion help.</p>

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

Alerts tab

On the **Alerts** tab, view alert messages that appear for certain events.

Note: The **Alerts** tab displays alert messages when a source schema change is detected. Messages are displayed 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:

- A **Custom** date range that consists of the beginning date and time and ending date and time that you select.

- **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.

The following table describes the columns of information that are displayed for each alert message:

Column	Description
Level	Severity level of the alert message.
Code	Alphanumeric code that identifies the alert type.
Details	Description of the event that raised the alert message.
Time	Date and time when the event occurred.

Note: You can configure alert notifications for application ingestion and database ingestion jobs from the **Alerts > Mass Ingestion Alerts** page in Operational Insights. Operational Insights then sends Mass Ingestion alert notifications to the users and user groups you select when ingestion jobs acquire certain statuses or detect a DDL change.

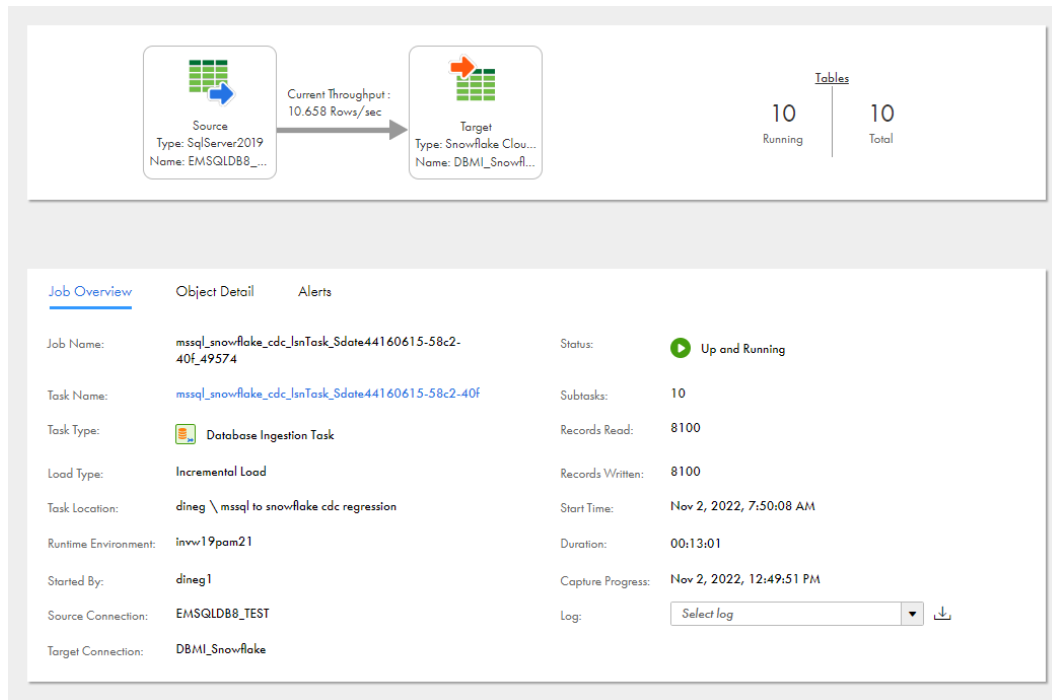
Database ingestion job details

For database ingestion jobs, you can view job-specific details on the **Job Overview**, **Object Detail**, and **Alerts** tabs. To access these tabs, drill down on a job from the **My Jobs** page in Mass Ingestion or from the **All Jobs** tab of the Mass Ingestion page in Operational Insights.

Job Overview tab

On the **Job Overview** tab, view detailed information for the entire job, including the name of the associated task, the load type, source and target connection names, current status, number of records read and written, start and end times, and run duration. For incremental load jobs and combined initial and incremental load jobs, you can also download the job log.

The following image shows the **Job Overview** tab for a database ingestion job:



Note: The diagram at the top of the page displays the calculated data throughput, in rows per second, if the job has successfully propagated 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.

The following table describes the **Job Overview** properties:

Property	Description
Job Name	The name of the job. Database ingestion job names have format <i>task name-job instance number</i> .
Task Name	The name of the associated ingestion task. You can click the task-name link to view or edit task details in Mass Ingestion, if necessary. If you edit the task, you must redeploy it for the updated task definition to be used for a job.
Task Type	The type of task, which is Database Ingestion Task .
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 Location	The project or project folder that contains the ingestion task definition.
Runtime Environment	The name of the runtime environment that the job uses to run.
Started By	The name of the user who started the job.

Property	Description
Source Connection	The name of the source connection.
Target Connection	The name of the target connection.
Status	<p>The status of the job, which can be one of the following values:</p> <ul style="list-style-type: none"> - 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 Mass Ingestion Databases (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.
Subtasks	The number of subtasks that the database ingestion job used to propagate data from source tables to the target. When a job runs, it uses a separate subtask to process each source table.
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 propagated to an Amazon S3, Apache Kafka, flat file, or Microsoft Azure Data Lake Storage target.</p> <p>For a Microsoft Azure Synapse Analytics target, the number of records written to the intermediate Microsoft Azure Data Lake Storage files.</p> <p>For a Snowflake target, the number of records written to the internal staging area that is created when the job runs.</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>
Start Time	The date and time when the job was deployed.
End Time	The date and time when the job ended because it completed processing, was stopped, or failed. This field is not displayed for running jobs
Duration	The amount of time, in the hh:mm:ss format, that the job ran before it ended.
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.

Property	Description
Capture Progress	<i>For experimental and developmental use.</i> For some incremental load and combined load sources, displays the date and UTC time in hh:mm:ss format that indicates the point in the source logs or change stream to which capture processing has progressed.
Log	<p>For incremental load jobs and combined initial and incremental load jobs, you can download the job execution log for the entire job run. Select one of the following log types:</p> <ul style="list-style-type: none"> - 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. <p>To download a log to your local system, click the Download icon.</p> <p>Note: For initial load jobs, you can get the job log for a specific source object from the Object Detail tab.</p>

Object Detail tab

On the **Object Detail** tab, view statistics and status information by source object from the last run of a database ingestion job.

The following image shows the **Object Detail** tab for a database ingestion job:

The screenshot displays the Informatica Operational Insights web interface. The left sidebar contains navigation links: Home, Alerts, All Infrastructure, Data Profiling, Application Info..., PowerCenter, Data Engineering..., Data Quality, Mass Ingestion, and Data Integration. The main content area shows a job overview with a source object (Type: Oracle Database, Name: Oracle_uf18...) and a target object (Type: Oracle Database, Name: RDS191_Krsh...). A table summary shows 1 Running and 1 Total table. Below this, the 'Object Detail' tab is active, showing a table with columns: Table Name, Target Object, Stage, Status, and Log. The table lists one item: AUSGQA_SRC_ALLNUM_BK, with a status of 'Running'. At the bottom, a summary table shows statistics for the job run:

Inserts	Updates	Deletes	LOBs	Unload Count
4	5	1	0	0

Note: This tab 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** properties that are displayed for each table, depending on the load type and status:

Column	Description
Table Name	<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 table.
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 an Amazon S3, Apache Kafka, flat-file, or Microsoft Azure Data Lake Storage target. For other load types, this information is available only at the job-level on the Job Overview tab.</p> <p>For a Microsoft Azure Synapse Analytics target, the number of records written to the intermediate Microsoft Azure Data Lake Storage files.</p> <p>For a Snowflake target, the number of records written to the internal staging area that is created when the job runs.</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.

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 Mass Ingestion Databases (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>For initial load jobs, you can download the job execution log for a source object. Select one of the following log types:</p> <ul style="list-style-type: none"> - Complete Log. The complete log for the object subtask from job execution. - Error. The log that contains error messages. This log type is available only for a Failed subtask. <p>For incremental load jobs, you can get the complete log for the entire job run from the Job Overview tab.</p> <p>For combined initial and incremental load jobs, you can download the Stage Log. This log covers the transition from initial to incremental loading for a source object.</p> <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>
Actions menu > Resync	<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 Mass Ingestion help.</p>

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

Alerts tab

On the **Alerts** tab, view alert messages that appear for certain events.

Note: The **Alerts** tab displays alert messages when a source schema change is detected. Messages are displayed 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:

- A **Custom** date range that consists of the beginning date and time and ending date and time that you select.
- **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.

The following table describes the columns of information that are displayed for each alert message:

Column	Description
Level	Severity level of the alert message.
Code	Alphanumeric code that identifies the alert type.
Details	Description of the event that raised the alert message.
Time	Date and time when the event occurred.

Note: You can configure alert notifications for application ingestion and database ingestion jobs from the **Alerts > Mass Ingestion Alerts** page in Operational Insights. Operational Insights then sends Mass Ingestion alert notifications to the users and user groups you select when ingestion jobs acquire certain statuses or detect a DDL change.

File ingestion job details

The job results for each file ingestion task instance display the status of the job, and success and error statistics.

To view detailed information about a file ingestion task, click the task name on the **My Jobs** page in Mass Ingestion or on the **All Jobs** tab of the Mass Ingestion page in Operational Insights.

You can download the job. The following image shows the details of a file ingestion job:

Job Properties		Results	
Task Name:	adlgen2_to_ftps_1SatJun122021094156	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 Message:	Job completed normally

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	
File6.txt	32	Success	ADLS Download	Jun 12, 2021, 10:31:29 AM	928	
File5.txt	35	Success	FTPS Upload	Jun 12, 2021, 10:31:28 AM	297	
File3.txt	45	Success	FTPS Upload	Jun 12, 2021, 10:31:27 AM	305	

Job Properties

The job properties for the file ingestion task instance display general properties about the instance.

The following table describes the job properties:

Property	Description
Task Name	The name of the associated ingestion task. You can click the task-name link to view or edit task details in Mass Ingestion.
Task Type	Task type. In this case, file ingestion 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 the file ingestion task instance display the status of the job and error statistics.

The job results include the following properties:

Property	Description
State	Job status. A job can have one of the following statuses: <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.
Session Log	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. Session log files are written to the following directory: <Secure Agent installation directory>/apps/Data_Integration_Server/logs
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 task has transferred with information about each file.

The File Events section is updated each time the file ingestion task transfers a file, and the state of the file updates 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 job to resume the transfer of the interrupted files. <p>Note: The status is applicable when the file ingestion task 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 task 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 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.

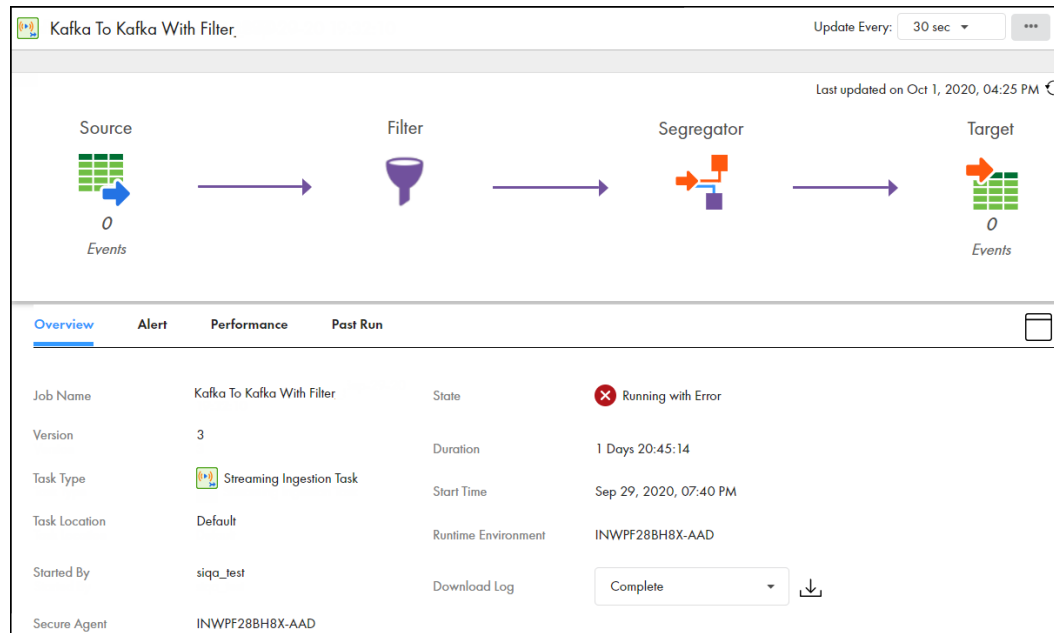
Streaming ingestion job details

To view detailed information about a streaming ingestion job, click the job name on the **My jobs** page in Mass Ingestion or on the **All Jobs** tab of the Mass Ingestion page in Operational Insights.

Overview tab

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

The following image shows the **Overview** tab for a streaming ingestion job:



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 task.
Task Location	The project or project folder that contains the streaming ingestion 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.
State	The state of the job. A job can have one of the following states: <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>Level of log that you want to download for a running job.</p> <p>You can download one of the following 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 state. - 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 job:

The screenshot shows the details of a streaming ingestion job titled "KafkaToFlatfileTargetWithFormatConverterWith...". The job is represented by a diagram with a "Source" icon (a green grid with a blue arrow) and a "FormatConverter" icon (a green square with a blue double arrow). Below the Source icon, the text "499021 Events" is displayed. Below the diagram, there are four tabs: "Overview", "Alert", "Performance", and "Past Run". The "Alert" tab is selected, showing a list of four alert messages. Each message starts with a yellow warning icon and contains the text: "KafkaToFlatfileTargetWithFormatConverterWithXML_FormatConverter : ConvertRecord[id=6ee4a30e-9ed9-4e00-87bf-d6cc0e0... Show M".

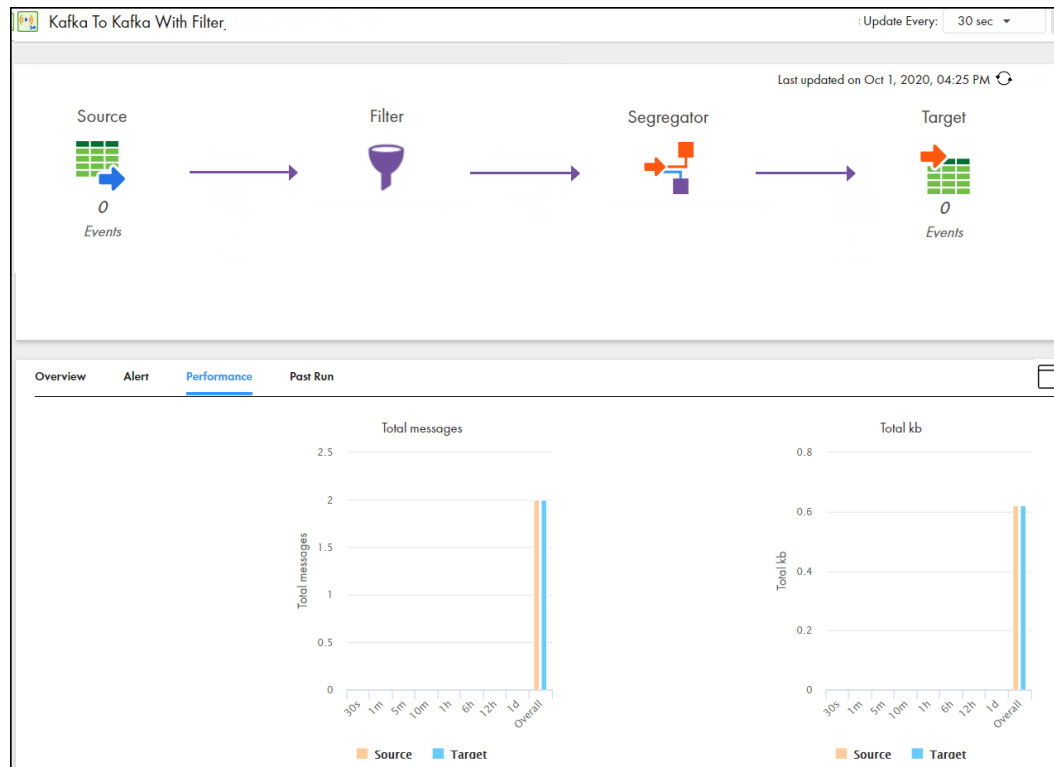
The following table describes the job alert properties:

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 for a streaming ingestion job:



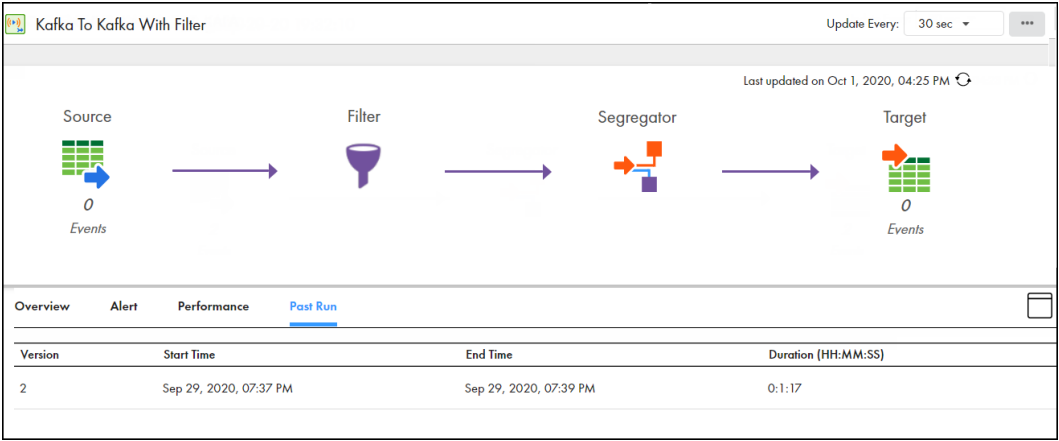
The following table describes the job performance properties:

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 related to the previous runs of a streaming ingestion job.

The following image shows the **Past Run** tab for a streaming ingestion job:



The following table describes the past run properties:

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 hh:mm:ss format.

INDEX

M

monitoring ingestion jobs
 database ingestion job details [15](#)
 Job Overview tab [15](#)
 job properties in job lists [9](#)
 monitoring all jobs from Operational Insights [6](#)
 monitoring My Jobs in Mass Ingestion [5](#)
 Object Details tab [15](#)

monitoring ingestion jobs (*continued*)
 overview [5](#)
 viewing job details [9](#)

S

streaming ingestion jobs [23](#)