Informatica® Intelligent Cloud Services
Spring 2020 April

Monitor

Table of Contents

Preface

Use *Monitor* to learn how to view and monitor jobs, imports, and exports that are running or have run in your Informatica Intelligent Cloud Services℠ organization. *Monitor* also contains information about monitoring at-scale clusters and source control actions and downloading log files to troubleshoot errors.

Informatica Resources

Informatica provides you with a range of product resources through the Informatica Network and other online portals. Use the resources to get the most from your Informatica products and solutions and to learn from other Informatica users and subject matter experts.

Informatica Documentation

Use the Informatica Documentation Portal to explore an extensive library of documentation for current and recent product releases. To explore the Documentation Portal, visit https://docs.informatica.com.

If you have questions, comments, or ideas about the product documentation, contact the Informatica Documentation team at infa_documentation@informatica.com.

Informatica Intelligent Cloud Services web site

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

Informatica Intelligent Cloud Services Communities

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

Access the Informatica Intelligent Cloud Services Community at:

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

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

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

Informatica Intelligent Cloud Services Marketplace

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

https://marketplace.informatica.com/

Data Integration connector documentation

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

Informatica Knowledge Base

Use the Informatica Knowledge Base to find product resources such as how-to articles, best practices, video tutorials, and answers to frequently asked questions.

To search the Knowledge Base, visit https://search.informatica.com. If you have questions, comments, or ideas about the Knowledge Base, contact the Informatica Knowledge Base team at KB_Feedback@informatica.com.

Informatica Intelligent Cloud Services Trust Center

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

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

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

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

Informatica Global Customer Support

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

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

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

CHAPTER 1

Monitoring jobs

You can monitor the jobs that are running or have run in your organization. A job is an instance of an asset such as a Data Integration mapping, task, or taskflow. Each time that you start the mapping, task, or taskflow, Informatica Intelligent Cloud Services creates a job to run it.

You can monitor jobs on the following pages:

All Jobs page

Lists all jobs that are running or that have run in your organization. For most job types, you can stop and restart the job on this page. You can also download log files for some job types.

To open the **All Jobs** page, in Monitor, select **All Jobs**.

Running Jobs page

Lists all jobs that are running or have completed within the last five minutes. For most job types, you can stop and restart the job on this page. You can also download log files for some job types.

To open the **Running Jobs** page, in Monitor, select **Running Jobs**.

My Jobs page

Lists all jobs that were started by the currently logged in user. For most job types, you can stop and restart the job on this page. You can also download log files for some job types.

To open the **My Jobs** page, in the service where you started the job, select **My Jobs**.

Job details page

Displays detailed information about a specific job. You can view details for jobs that have completed or failed. For most job types, you can restart the job from this page. You can also download log files for some job types.

To view the details for a job, click the job name on the **All Jobs**, **Running Jobs**, or **My Jobs** page.

Monitoring all jobs

You can monitor all jobs in your organization on the **All Jobs** page in Monitor. The **All Jobs** page lists the jobs that are currently running and the jobs that have completed.

Use the **All Jobs** page for failure analysis and debugging of the jobs in your organization.

To monitor jobs on the **All Jobs** page, you need one of the following roles:

- Administrator
- Designer

- Monitor

The following image shows the **All Jobs** page:
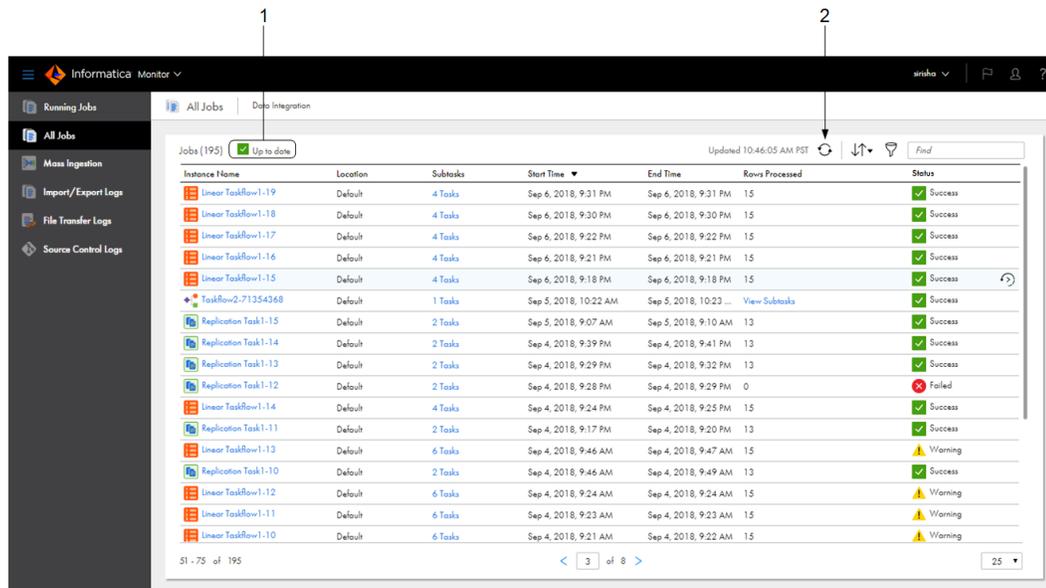
1. Status message that indicates whether information on the page is up-to-date or needs to be refreshed
2. Refresh icon

The **All Jobs** page lists the jobs that were run within the last three days, plus the 1000 most recent jobs that are more than three days old.

To ensure that the information on this page is current, Monitor polls the Informatica Intelligent Cloud Services repository every five seconds. Information can become out-of-date when a job status changes or when a user starts a job.

The status message at the top of the page indicates whether the information on the page is up-to-date. If the information is out-of-date, the status message displays "Updates Available." To refresh the page, click the "Updates Available" message or the **Refresh** icon.

When a job completes, you can drill down on the job to view the job details. To drill down on a job, click the instance name.

Monitoring running jobs

You can monitor all running jobs on the **Running Jobs** page in Monitor. The **Running Jobs** page lists the jobs that are starting, queued, running, and suspended. The page also lists the jobs that have completed within the last five minutes.

Use the **Running Jobs** page for live monitoring of the jobs that are running in your organization. When you view running jobs, job properties such as end time, rows processed, and status are continuously updated.

To monitor jobs on the **Running Jobs** page, you need one of the following roles:

- Administrator
- Designer

- Monitor

The following image shows the **Running Jobs** page:

If your organization is running more than 200 jobs, the jobs are listed on multiple pages.

When a job completes, you can drill down on the job to view the job details. To drill down on a job, click the instance name.

Monitoring your jobs

You can monitor the jobs that you started on the **My Jobs** page within the service where you started the job. The **My Jobs** page lists your jobs that are currently running and your jobs that have completed.

You can view the **My Jobs** page in the following services:

- Data Integration
- Data Profiling
- Mass Ingestion

The following image shows the **My Jobs** page in Data Integration:
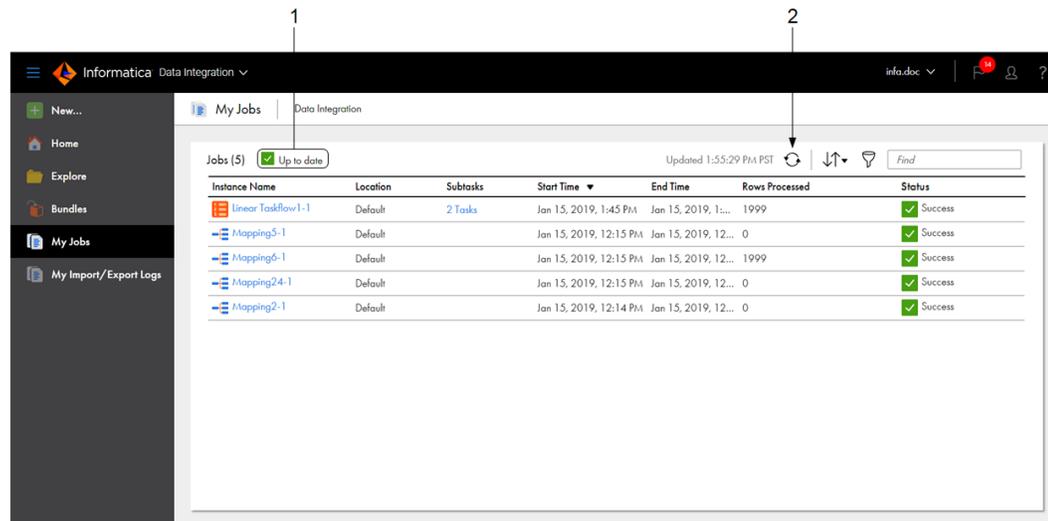
1. Status message that indicates whether information on the page is up-to-date or needs to be refreshed.
2. Refresh icon

The **My Jobs** page lists the jobs that you ran within the last three days, plus your 1000 most recent jobs that are more than three days old.

To ensure that the information on this page is current, Monitor polls the Informatica Intelligent Cloud Services repository every five seconds. Information can become out-of-date when a job status changes or when a user starts a job.

The status message at the top of the page indicates whether the information on the page is up-to-date. If the information is out-of-date, the status message displays "Updates Available." To refresh the page, click the "Updates Available" message or the **Refresh** icon.

When a job completes, you can drill down on the job to view the job details. To drill down on a job, click the instance name.

Monitoring subtasks

When you monitor a replication task, taskflow instance, or an Azure data sync task, you can also monitor the subtasks. You can monitor a subtask from the **All Jobs** or **Running Jobs** page in Monitor or from the **My Jobs** page in the service where you started the job.

To monitor subtasks, open the **All Jobs**, **Running Jobs**, or **My Jobs** page, and then click the number of subtasks in the Subtasks column. For non-linear taskflows, you can also click **View Subtasks** in the Rows Processed column.

The following types of assets have subtasks:

- Replication tasks
- Advanced taskflows
- Linear taskflows
- Azure data sync tasks

- Data Profiling tasks

The following image shows the **My Jobs** page in Data Integration when you view subtasks for a linear taskflow:
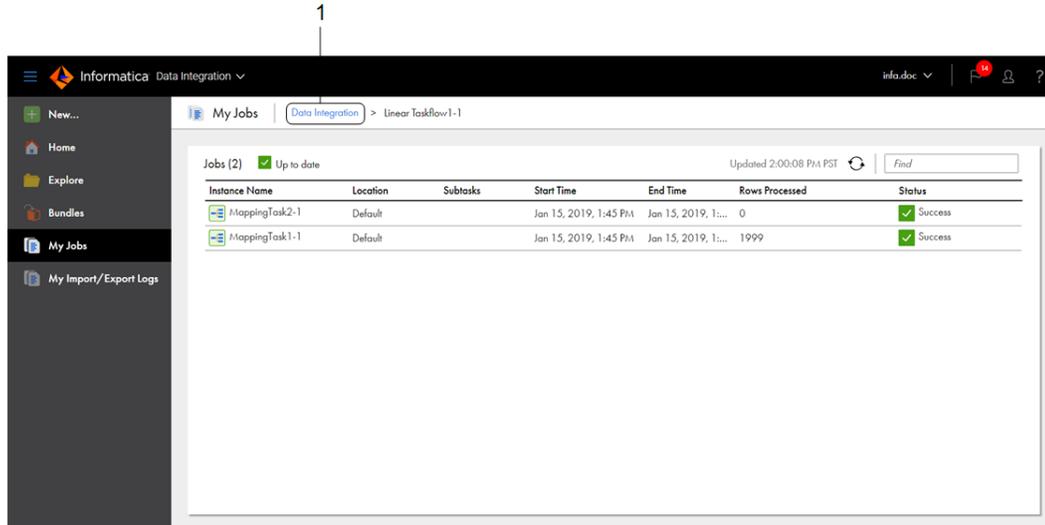
1. Click the service name to close the subtasks view.

To close the subtasks view and return to the **My Jobs** page, click the service name in the menu at the top of the page. For example, to close the subtasks view for a linear taskflow, click **Data Integration** in the menu at the top of the page to the left of the parent task or taskflow instance name.

Job properties

The **All Jobs**, **Running Jobs**, and **My Jobs** pages display the job properties such as the name, start time, and status. You can right-click the column heading area to display or hide specific properties.

The following image shows the default properties that are displayed for mapping, task, and taskflow instances on the **All Jobs** page in Data Integration:
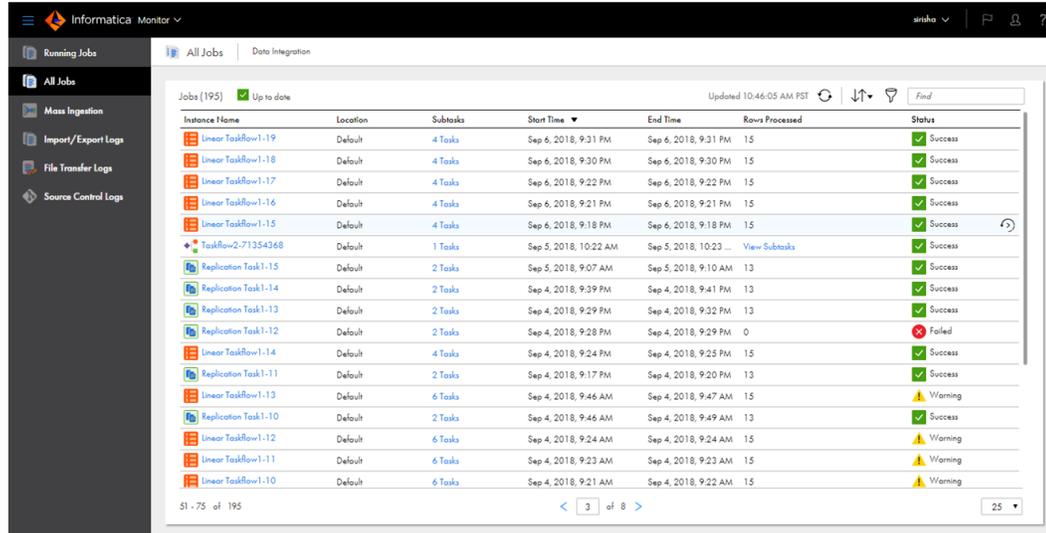
Each Informatica Intelligent Cloud Services job is named `<asset name>-<instance number>`. For example, the first time you run mapping m_LoadCustOrders, the job is named m_LoadCustOrders-1. The second time you run the mapping, the job is named m_LoadCustOrders-2.

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

Property	Description
Instance Name	Name of the job in the following format: <asset name>-<instance number> For completed and failed jobs, you can click the instance name to view detailed information about the job.
Location	Project and folder path where the asset exists.
Subtasks	Displays the number of subtasks, when applicable. To view job properties for the subtasks, click the entry in this column.
Start Time	Date and time that the job was started.
End Time	Date and time that the job completed or stopped. Does not apply to running jobs.

Property	Description
Rows Processed	Total number of rows that the job has currently processed. This value includes the number of rows successfully written to the target as well as the number of error rows for each source, target, and transformation in the task. If the job you are viewing is a non-linear taskflow instance, this field displays **View Subtasks**. Click **View Subtasks** to monitor the subtasks for the taskflow. Does not apply to elastic mappings.
Status	Job status. A job can have one of the following statuses: - Queued. The job is queued on a Secure Agent, but it has not started yet. Applies to subtasks of replication task and taskflow instances. - Starting. The job is starting. - Running. The job is still running. - Success. The job completed successfully. - Stopped. The parent job has stopped running, so the subtask cannot start. Applies to subtasks of replication task instances. - Suspended. The job is paused. Applies to taskflow instances. - Warning. The job completed with errors. Applies to mapping and task instances. - Aborted. The job was aborted. Applies to mass ingestion task instances. - Failed. The job did not complete because it encountered errors. **Tip:** If the job status is Failed, you can hover the cursor over the job status to view and copy the error message.

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

Property	Description
Asset Name	Name of the asset that is associated with the job. For example, if the job is a mapping instance, this column displays the name of the associated mapping.
Asset Type	Type of asset that is associated with the job.
Runtime Environment	Runtime environment in which the job ran.
Duration	Amount of time the job ran before it completed or was stopped.
Success Rows	Number of rows successfully written to the target. Does not apply to mass ingestion tasks or mapping tasks that are based on elastic mappings.
Failure Rows	Number of rows that were not written to the target. Does not apply to mass ingestion tasks or mapping tasks that are based on elastic mappings.
Started By	Name of the user or schedule that started the job.
Error Message	Error message, if any, that is associated with the job.

Customizing the jobs pages

You can decide which properties to display on the **All Jobs**, **Running Jobs**, and **My Jobs** pages. You can also sort and filter jobs, drill down on a job to view detailed information about the job, and view properties for subtasks.

The following image shows the properties that are displayed by default on the **All Jobs** page:
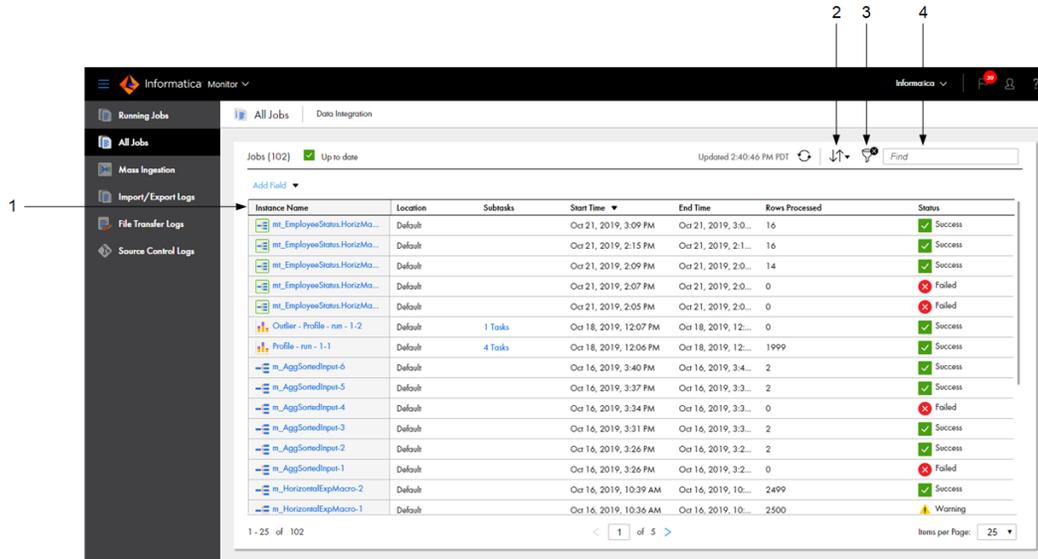
1. Column heading area. Right-click this area to add and remove columns.
2. Sort icon
3. Filter icon
4. Find field

You can customize the jobs pages in the following ways:

Display, hide, or rearrange job properties.

To display or hide specific properties, right-click the column heading area and check or uncheck the properties.

To rearrange the columns, click a column heading and drag it to a different location.

Sort jobs.

To sort the displayed jobs, click the column heading for the property that you want to sort by. For example, to list the most recently completed jobs first, click the End Time column. The arrow in the column heading indicates the sort order, either ascending or descending. To reverse the sort order, click the column heading again.

You can also sort jobs by clicking the **Sort** icon and selecting the column name.

Find jobs.

To find specific jobs on the **All Jobs**, **Running Jobs**, and **My Jobs** pages, use the find field. You can search for jobs by entering a full or partial job name or error message in the find field.

Filter jobs.

To filter the jobs that appear on the **All Jobs** and **My Jobs** pages, click the **Filter** icon. Use filters to find specific jobs. You can specify keywords and partial strings in your filters.

To specify a filter, click **Add Field**, select the property to filter by, and then enter the property value. For example, to find jobs with names that contain the string "Sales," select "Asset Name" as the filter field, and enter "Sales" as the value. Or, to find jobs with errors that contain the string "Hosted Agent," select "Error Message" as the filter field, and enter "Hosted Agent" as the value.

You can specify multiple filters. For example, to find Data Integration failed mappings, select the following filter fields and values:

- Asset Type: Mapping
- Status: Failed

The following image shows the **All Jobs** page with the filters applied:
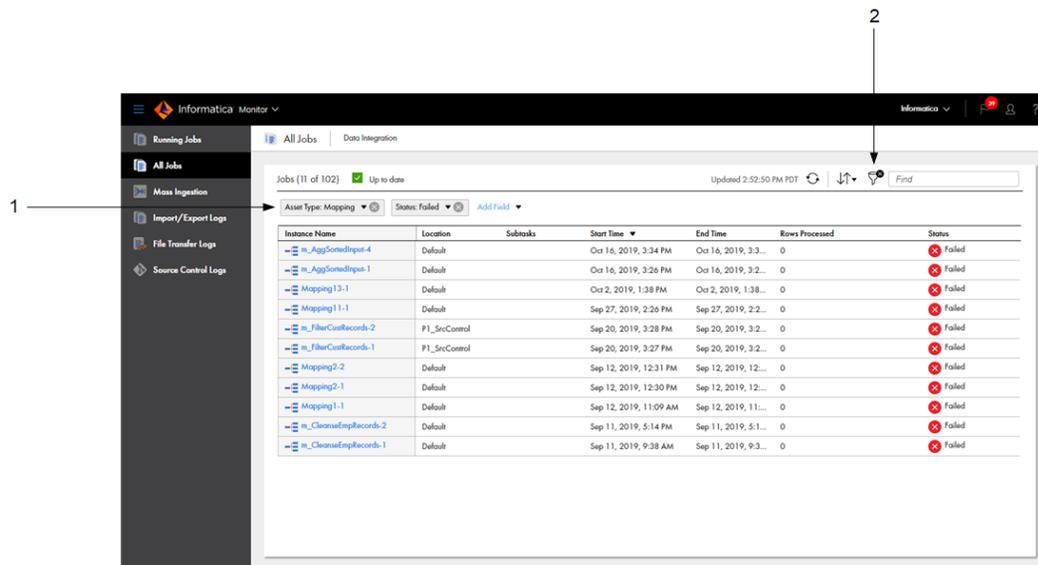
1. Applied filters
2. Remove Filter icon

To remove all applied filters, click the **Remove Filter** icon.

Viewing details for a specific job

You can drill down on any job on the **All Jobs**, **Running Jobs**, or **My Jobs** page to display detailed information about the job.

To drill down and view details about a specific job, click the job name.

To view details for jobs with subtasks, open the subtasks view, and click the subtask name.

The following image shows the details for a linear taskflow instance after you drill down:

The details displayed for the job vary based on the job type.

CHAPTER 2

Monitoring Data Integration jobs

You can monitor the Data Integration jobs that are running or have run. A job is an instance of a mapping, task, or taskflow. Each time that you start a mapping, task, or taskflow, Data Integration creates a job to run it.

You can monitor the following types of Data Integration jobs:

- Mappings
- Mapping data preview jobs
- Mapping tasks
- Elastic mappings
- Synchronization tasks
- Masking tasks
- Replication tasks
- PowerCenter tasks
- Mass ingestion tasks
- Linear taskflows
- Advanced taskflows

Monitoring mappings and mapping tasks

To view detailed information about a specific mapping instance, task instance, or mapping data preview job, click the job name on the **All Jobs**, **Running Jobs**, or **My Jobs** page.

The following image shows the job details for a mapping task:
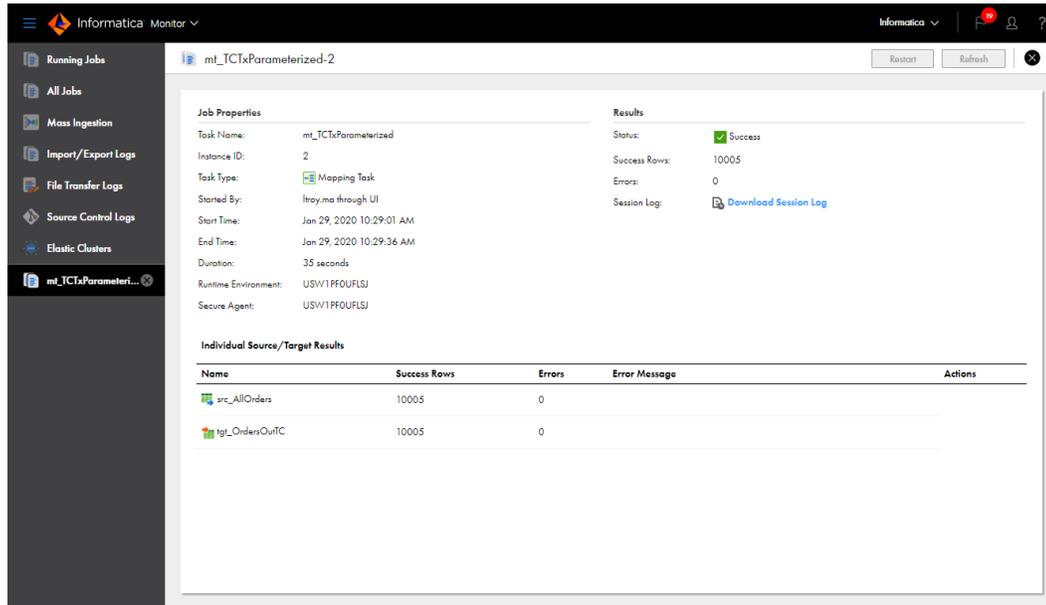
You view different details in each panel, and the details change based on the type of task that you run.

Job properties

The job properties for each mapping instance, task instance, or mapping data preview job display general properties about the instance.

The job properties include the following properties:

Property	Description
Task name	Name of the task.
Instance ID	Instance number for the task. For example, if you are looking at the third run of the task, this field displays "3."
Task type	Task type, for example mapping task.
Started by	Name of the user or schedule that started the job.
Start time	Date and time that the job was started.
End time	Date and time that the job completed or stopped.
Duration	Amount of time the job ran before it completed or was stopped.

Property	Description
Runtime environment	Runtime environment in which the job ran.
Secure Agent	Secure Agent that ran the job.

Job results

The job results for each mapping instance, task instance, or mapping data preview job display the status of the job and success and error statistics.

The job results include the following properties:

Property	Description
Status	Job status. A job can have one of the following statuses: - Starting. The job is starting. - Running. The job is still running. - Success. The job completed successfully. - Warning. The job completed with errors. - Failed. The job did not complete because it encountered errors.
Success Rows	Number of rows successfully written to the target.
Errors	Total number of source error rows, target error rows, and transformation errors.
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 want to retain session logs for more than 10 runs, you can configure the **Maximum Number of Log Files** property in the task wizard. Session log files are written to the following directory: `<Secure Agent installation directory>/apps/Data_Integration_Server/logs`
Error Message	Error message, if any, that is associated with the job.

Individual source and target results

Individual source and target results are displayed for mapping task instances and mapping data preview jobs that do not apply to elastic mappings. Individual source and target results display details for individual sources and targets.

The individual source and target details include the following properties for each source and target:

Property	Description
Name	Name of the source or target transformation.
Success rows	Number of rows successfully written to the target.
Errors	Total number of source error rows, target error rows, and transformation errors.

Property	Description
Error message	Error message, if any, that is associated with the job.
Actions	Actions that you can take regarding the job.

If the task includes connection or object parameters that were overrun at runtime, the job details display the location of the parameter file and the value of each parameter after the task has completed.

If the task includes in-out parameters, the job details also display the value for each parameter after the task has completed.

Monitoring elastic mappings and tasks

To view detailed information about a specific elastic job, such as an elastic mapping or an associated mapping task, click the job name on the **All Jobs**, **Running Jobs**, or **My Jobs** page.

Note: When you monitor an elastic job while the job is running, you must refresh the page to view updates to the job properties, job results, and Spark task details.

Job properties

The job properties for each elastic mapping or task instance display general properties about the instance.

The following table describe the job properties for elastic mappings and tasks:

Property	Description
Task Name	Name of the task.
Instance ID	Instance number for the task. For example, if you are looking at the third run of the task, this field displays "3."
Task Type	Task type, in this case, mapping task.
Started By	Name of the user or schedule that started the job.
Start Time	Date and time that the job was started.
End Time	Date and time that the job completed or stopped.
Duration	Amount of time the job ran before it completed or was stopped.
Runtime Environment	Runtime environment in which the job ran.
Elastic configuration	Elastic configuration that was used to create the elastic cluster.
Cluster	Elastic cluster where the elastic job runs. You can click the cluster name to navigate directly to the monitoring details for the cluster.

Job results

The job results for each elastic mapping or task instance display the status of the job and success and error statistics.

The job results include the following properties:

Property	Description
Status	Job status. A job can have one of the following statuses: - Starting. The job is starting. - Running. The job is still running. - Success. The job completed successfully. - Failed. The job did not complete because it encountered errors. If the elastic cluster is not running when you run an elastic job, the job waits for the cluster to start. During this time the job status is Starting. If the Secure Agent fails while the elastic job is running, the status of the job continues to display Running. You must cancel the job and run the job again.
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 want to retain session logs for more than 10 runs, you can configure the **Maximum Number of Log Files** property in the task wizard. Session log files are written to the following directory: `<Secure Agent installation directory>/apps/Data_Integration_Server/logs`
Execution Plan	Allows you to download the Spark execution plan for a mapping task that is based on an elastic mapping. The execution plan shows the runtime Scala code that the Serverless Spark engine uses to run the data logic in the mapping. You can use the Scala code to debug issues in the elastic mapping.
Error Message	Error message, if any, that is associated with the job.

Job results for tuning

If you tune a mapping task that is based on an elastic mapping, the job results show the tuning job status and a link to the subtasks.

The job results include the following properties:

Property	Description
Status	Job status. A job can have one of the following statuses: - Running. The job is still running. - Success. The job completed successfully. - Stopped. The job was stopped. - Failed. The job did not complete because it encountered errors. If the Secure Agent fails while one of the subtasks is running, the statuses of the subtask and the tuning job display Running. You must stop the tuning job and configure tuning from the mapping task details again.
Subtasks	Number of subtasks that are part of the tuning job. Each subtask represents a run of the mapping task. If a link is available, click the link to monitor each mapping task.
Error Message	Error message, if any, that is associated with the job.

Spark task details

For mapping tasks that are based on elastic mappings, the mappings are translated into Spark tasks that process the data logic simultaneously. You can view details for each Spark task.

Each Spark task includes the following details:

Property	Description
Name	Name of the Spark task.
Status	Status of the Spark task. The Spark task can have one of the following statuses: - Running. The task is running. - Succeeded. The task completed successfully. - Failed. The task did not complete because it encountered errors. - Stopped. The task was stopped. - Unknown. The status of the task is unknown. If the Secure Agent fails while the elastic job is running, the status of the Spark tasks continues to display Running. You must cancel the job and run the job again.
Start time	Date and time when the Spark task started.
End time	Date and time when the Spark task ended.

Property	Description
Error message	Error encountered when running the Spark task, if any.

Actions	Actions that you can take regarding the Spark task. You can perform the following actions:
	- Download the Spark driver and agent job logs.
	- View the advanced log location.
	   The advanced log location is the log location that is configured in the elastic configuration for the elastic cluster.
	   You can navigate to the advanced log location to view copies of the agent job log, Spark driver log, and Spark executor logs.
	- Download advanced logs. If you use a serverless runtime environment, you can download the agent job log, Spark driver log, and Spark executor logs.

Monitoring synchronization tasks

To view detailed information about a specific synchronization task instance, click the job name on the **All Jobs**, **Running Jobs**, or **My Jobs** page.

Job properties

The job properties for the synchronization task display general properties about the instance.

The job properties include the following properties:

Property	Description
Task Name	Name of the task.
Instance ID	Instance number for the task. For example, if you are looking at the third run of the task, this field displays "3."
Task Type	Task type, in this instance, synchronization task.
Started By	Name of the user or schedule that started the job.
Start Time	Date and time that the job was started.
End Time	Date and time that the job completed or stopped.
Duration	Amount of time the job ran before it completed or was stopped.
Runtime Environment	Runtime environment in which the job ran.
Secure Agent	Secure Agent that ran the job.

Job results

The job results for each synchronization task display the status of the job and success and error statistics.

The job results include the following properties:

Property	Description
Status	Job status. A job can have one of the following statuses: - Starting. The job is starting. - Running. The job is still running. - Success. The job completed successfully. - Warning. The job completed with errors. - Failed. The job did not complete because it encountered errors.
Success Rows	Number of rows successfully written to the target.
Errors	Total number of source error rows, target error rows, and transformation errors.
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 want to retain session logs for more than 10 runs, you can configure the **Maximum Number of Log Files** property in the task wizard. Session log files are written to the following directory: `<Secure Agent installation directory>/apps/Data_Integration_Server/logs`
Error Rows File	Allows you to preview the error rows file for instances that have error rows. Click **Preview Error Rows File** to download the first 25 error rows and first 50 columns as a CSV file.
Error Message	Error message, if any, that is associated with the job.

Monitoring masking tasks

To view detailed information about a specific masking task instance, click the job name on the **All Jobs**, **Running Jobs**, or **My Jobs** page.

Job properties

The job properties for each masking task instance display general properties about the instance.

The job properties include the following properties:

Property	Description
Task Name	Name of the task.
Instance ID	Instance number for the task. For example, if you are looking at the third run of the task, this field displays "3."
Task Type	Task type, in this instance, masking task.

Property	Description
Started By	Name of the user or schedule that started the job.
Start Time	Date and time that the job was started.
End Time	Date and time that the job completed or stopped.
Duration	Amount of time the job ran before it completed or was stopped.
Runtime Environment	Runtime environment in which the job ran.
Secure Agent	Secure Agent that ran the job.

Job results

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

The job results include the following properties:

Property	Description
Status	Job status. A job can have one of the following statuses: - Starting. The job is starting. - Running. The job is still running. - Success. The job completed successfully. - Warning. The job completed with errors. - Failed. The job did not complete because it encountered errors.
Success Rows	Number of rows successfully written to the target.
Errors	Total number of source error rows, target error rows, and transformation errors.
Error Rows File	Allows you to preview the error rows file. Click **Preview Error Rows File** to download the first 25 error rows and first 50 columns as a CSV file.
Subset Rows	Number data subset rows loaded to the target.
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`
Error Message	Error message, if any, that is associated with the job.

Activity log details

The activity log displays results for individual subtasks.

The activity log details include the following properties for each subtask:

Property	Description
Name	Name of the subtask
End time	Date and time the subtask completed. You can also download the session log for the subtask.
Status	Status of the subtask.
Success rows	Number of rows successfully written to the target.
Error rows	Number of rows that were not written to the target
Error message	Error message, if any, associated with the job.
Recommendation	Recommended action you can take regarding the job.

Monitoring replication tasks

To view detailed information about a specific replication task instance, click the job name on the **All Jobs**, **Running Jobs**, or **My Jobs** page.

Job properties

The job properties for each replication task display general properties about the instance.

The job properties include the following properties:

Property	Description
Task Name	Name of the task.
Instance ID	Instance number for the task. For example, if you are looking at the third run of the task, this field displays "3."
Task Type	Task type, in this instance, replication task.
Started By	Name of the user or schedule that started the job.
Start Time	Date and time that the job was started.
End Time	Date and time that the job completed or stopped.
Duration	Amount of time the job ran before it completed or was stopped.

Property	Description
Runtime Environment	Runtime environment in which the job ran.
Secure Agent	Secure Agent that ran the job.

Job results

The job results for each replication task display the status of the job and success and error statistics.

The job results include the following properties:

Property	Description
Status	Job status. A job can have one of the following statuses: - Queued. The job is queued on a Secure Agent, but it has not started yet. - Starting. The job is starting. - Running. The job is still running. - Success. The job completed successfully. - Stopped. The parent job has stopped running, so the subtask cannot start. - Warning. The job completed with errors. - Failed. The job did not complete because it encountered errors.
Success Rows	Number of rows successfully written to the target.
Errors	Total number of source error rows, target error rows, and transformation errors.
Error Message	Error message, if any, that is associated with the job.

Individual object results

The individual object results display results for individual target objects. You can also download the session logs for individual subtasks.

The job details include the following properties for each target object:

Property	Description
Object name	Name of the target object.
End time	Date and time that the job completed or stopped.
Status	Job status, either running, success, warning, or failed.
Success rows	Number of rows successfully written to the target.
Errors	Total number of source error rows, target error rows, and transformation errors.
Error message	Error encountered when writing to the target, if any.
Recommendation	Recommended action for errors.

Monitoring PowerCenter tasks

To view detailed information about a specific PowerCenter task, click the job name on the **All Jobs**, **Running Jobs**, or **My Jobs** page.

Job properties

The job properties for each PowerCenter task instance display general properties about the instance.

The job properties include the following properties:

Property	Description
Task Name	Name of the task.
Instance ID	Instance number for the task. For example, if you are looking at the third run of the task, this field displays "3."
Task Type	Task type, in this instance, PowerCenter.
Started By	Name of the user or schedule that started the job.
Start Time	Date and time that the job was started.
End Time	Date and time that the job completed or stopped.
Duration	Amount of time the job ran before it completed or was stopped.
Runtime Environment	Runtime environment in which the job ran.
Secure Agent	Secure Agent that ran the job.

Job results

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

The job results include the following properties:

Property	Description
Status	Job status. A job can have one of the following statuses: - Starting. The job is starting. - Running. The job is still running. - Success. The job completed successfully. - Warning. The job completed with errors. - Failed. The job did not complete because it encountered errors.
Success Rows	Number of rows successfully written to the target.
Errors	Total number of source error rows, target error rows, and transformation errors.

Property	Description
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`
Error Message	Error message, if any, that is associated with the job.

Monitoring mass ingestion tasks

To view detailed information about a specific mass ingestion task instance, click the job name on the **All Jobs**, **Running Jobs**, or **My Jobs** page.

You cannot stop or restart a mass ingestion task in Monitor.

Job properties

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

The job properties include the following properties:

Property	Description
Task Name	Name of the task.
Task Type	Task type, in this case, mass ingestion task.
Started By	Name of the user or schedule that started the job.
Start Time	Date and time that the job was started.
End Time	Date and time that the job completed or stopped.
Duration	Amount of time the job ran before it completed or was stopped.

Job results

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

The job results include the following properties:

Property	Description
Status	Job status. A job can have one of the following statuses: - Starting. The job is starting. - Running. The job is still running. - Success. The job completed successfully. - Warning. The job completed with errors. - 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

The File Events section shows the total number of files that the mass ingestion task has transferred, with information about each file.

The File Events section is updated each time the mass ingestion task transfers a file, and the state of the file updates throughout the file transfer process. You can track the progress of the transfer of a file 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.

Property	Description
State	The state of the file transfer. A file can have one of the following states: - 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 mass ingestion task to resume the transfer of the interrupted files. This state is applicable when the mass ingestion task transfers file from the advanced FTP, advanced SFTP, or advanced FTPS sources. - In Doubt. The previous task instance encountered errors while transferring the file. Applicable for tasks where the source is configured to skip duplicate files. You can monitor the State property to track the progress of the file transfer of each file.
Transfer Type	The type of file transfer. A file can have one of the following transfer types: - *<source_name>* Download. The file is downloaded from source. *<source_name>* is the name of the source. - *<target_name>* Upload. The file is uploaded to the target. *<target_name>* is the name of the target. - Copy from Source. The mass 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	Applicable for file events in Failed status. The message includes the reason for failure of the event based on the file transfer type.

Monitoring taskflows

To view detailed information about a taskflow instance, click the job name on the **Jobs** or **My Jobs** page.

The following image shows the job details for a taskflow instance:
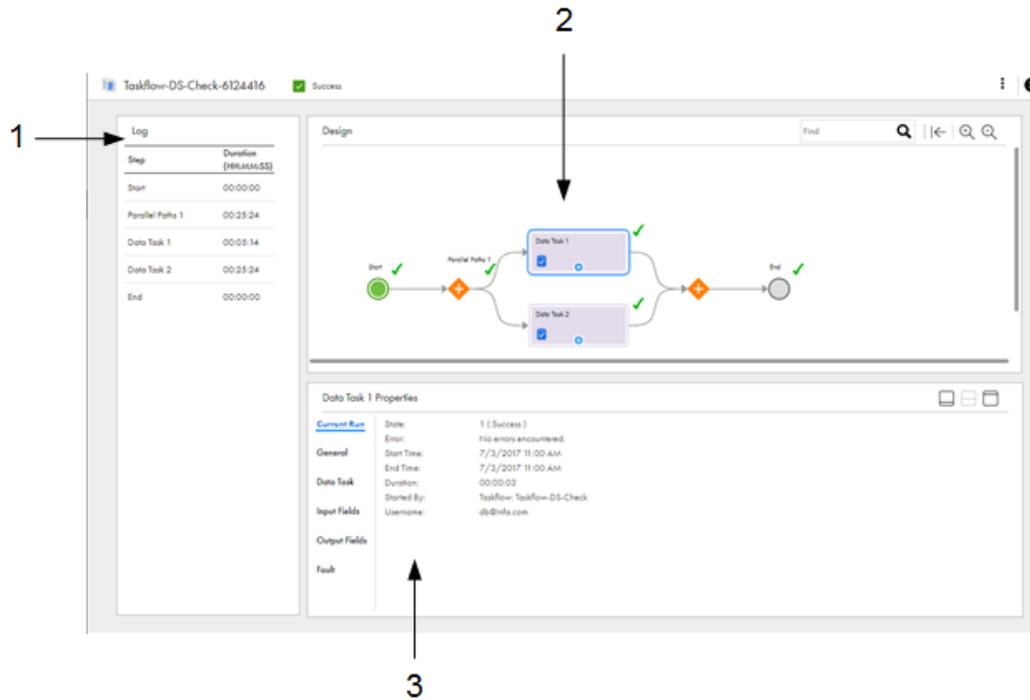
1. Log panel. A chronological list of the activities that the taskflow has run.
2. Design panel. A canvas view of the taskflow.
3. Properties panel. The properties of the taskflow element that you select on the canvas. Click the Start step to view properties specific to the current taskflow instance. Click a Data Task step to see properties specific to the task.

Start step properties

The Start step properties display information about the current run of a taskflow.

When you click the Start step, the properties panel displays the following information:

Property	Description
Instance ID	A unique number that identifies the current taskflow instance.
State	The state that the taskflow instance is in. Taskflows can be in the Running, Failed, Success or Suspended states.
Start Time	The date and time at which the taskflow instance was started.
End Time	The date and time at which the taskflow instance ended. You see an End Time value if a taskflow is in the Success or Failed states.
Duration	The amount of time the taskflow instance ran before it completed. You see a Duration value if a taskflow is in the Success or Failed states.

Property	Description
Username	The name of the user who ran this instance of the taskflow.
Runtime Environment	The runtime environment in which the taskflow ran. You can run taskflows only in the Cloud environment.

Data Task properties

The Data Task step properties display details about the current run of the data task of a taskflow.

When you click a data task, the properties panel displays the following information:

Property	Description
Current Run	You see the following properties: - State. Indicates the state of the data integration task. A data integration task can be in the **0 (Faulted)**, **1 (Success)**, or **2 (Warning)** states. - Error. If the data integration task is on the Faulted or Warning states, the error message appears here. - Start Time. The date and time at which the data integration task started. - End Time. The date and time at which the data integration task ended or was suspended. - Duration. The amount of time the data integration task ran before it completed or was suspended. - Started By. The name of the taskflow that called the data integration task. - Username. The name of the user who ran this instance of the taskflow.
General	The name of the Data Task step.
Data Task	You see the following properties: - Name. The name of the data integration task. - Type. Whether the data integration task is a mapping task or a synchronization task.
Input Fields	Properties specific to the data integration task you added to a Data Task step. You see the following properties: - Wait for Task to Complete. Indicates if the Data Task step waits for the data integration task to complete. This value is always **Yes**. - Max Wait. The maximum time in seconds that the Data Task step waits for the data integration task to complete. This value is always **604800** seconds, which is 7 days. - Task Name. The name of the mapping task or synchronization task. - GUID. The unique identifying number of the data integration task. - Task Type. Indicates whether the data integration task is a mapping task or a synchronization task.
Output Fields	You see Output Field properties if the **State** of the data integration task is **1 (Success)**. Properties specific to the completed data integration task appear here. **Note:** If the data task is a mapping task that is based on an elastic mapping, the Output Field properties include success target rows and error target rows, but they do not include success source rows and error source rows.
Faults	If the data integration task faulted, you see the following properties: - Code. Indicates whether the fault is an error or a warning. - Reason. Details about the error or fault.

Notification Task properties

The Notification Task step properties display details about the current run of the Notification Task of a taskflow.

When you click a Notification Task, the **Properties** panel displays the following information:

Property	Description
General	The name of the Notification Task step.
Details	Displays the following input fields of the Notification Task: - Email To. The primary recipients for the email notification. - Email Cc. The recipients who need to be sent a copy of the email notification. - Email Subject. A short and descriptive subject that introduces the email. - Email Content Type. The type of formatting that the email content uses. - Email Body. The content that the email contains.

Command Task step properties

The Command Task step properties display details of the current run of the command task in a taskflow. When you click a command task in the taskflow, the **Properties** panel displays the following properties:

Property	Description
General	Displays the name of the command task.
Input Fields	Displays the value of the runtime environment, script file name, input arguments, and work directory that you configure for the command task.
Output Fields	Displays the value of the run ID, start and end time, exit code, execution status, and standard error. The output fields are displayed only if the command task executes successfully.
Fault	Displays the error code and reason for the error.

File Watch Task step properties

The File Watch Task step properties display details about the current run of the file watch task of a taskflow.

When you click a File Watch Task step, the **Properties** panel displays the following information:

Property	Description
Current Run	Displays the following details: - Status. Status of the file watch task. - Started By. Name of the taskflow that called the file watch task.
General	Displays the name of the File Watch Task step.
File Watch Task	Displays the name and type of the file listener that the File Watch Task step contains. The type is always set to **MI_FILE_LISTENER**.

Property	Description
Input Fields	Displays the following input fields: - Wait for Task to Complete. Indicates whether the File Watch Task step waits for the file listener to complete. This value is always set to **true**. - Max Wait. Maximum time in seconds that the File Watch Task step waits for the file listener to complete. This value is always set to **604800** seconds, which is 7 days. - Task Name. Name of the file listener. - GUID. Unique identification number for the file listener. - Task Type. Indicates the type of file listener that the File Watch Task step contains. This value is always set to **MI_FILE_LISTENER**. - serviceProviderId. Service provider ID for the file listener. This value is always set to **filelistener**. - serviceProviderContextPath. Service provider context path for the file listener. This value is always set to **filelistener/**.

Output Fields	Displays the following output fields if the File Watch Task step completed successfully:
	
	- id. ID for the File Watch Task step run.
	- monitorJobId. Job ID that you can use to check the status of the file listener in the **File Transfer Logs** tab of the Monitor service.
	
	  Append the monitorJobId to the file listener name to find out the instance name in the **File Transfer Logs** tab of the Monitor service.
	
	  For example, if the monitorJobId is **7500** and the name of the file listener is **FL_Arrive**, the instance name that you need to look for in the **File Transfer Logs** tab of the Monitor service would be **FL_Arrive-7500**.
	
	- notificationTime. Timestamp when the File Watch Task step completed. The timestamp includes the date and time up to seconds.
	- status. Status of the file watch task. This value is always set to **Completed**.
	- fileEvents. Provides a list of files along with the following details:
	  - Event ID. Auto-generated ID for the event.
	  - Event type. Indicates the nature of the file event. You see one of the following values:
	    - ARRIVED
	    - UPDATED
	    - DELETED
	  - File last modified. Timestamp when the file was last modified. The timestamp includes the date and time up to seconds.
	  - File name. Name of the file that arrived, was updated, or deleted as part of the event.
	  - File path. Name of the file and the path where the file exists.
	  - File size. Size of the file in KB.
	- Task Id. This value is always blank.
	- Run Id. This field has the same value as the **monitorJobId** field.

Fault	If the file watch task faulted, you see the following properties: - Code. Indicates whether the fault is an error or a warning. - Reason. Cause of the error or fault. - Details. Details about the error or fault.

Consider the following rules and guidelines when you monitor the File Watch Task step:

- The **All Jobs** page does not display the file watch task as a subtask. It always displays the number of rows processed as **0** if the taskflow includes a File Watch Task step.

- If the file watch task faulted and the **On Error** property was set to **Suspend Taskflow**, no information is displayed in the **Properties** panel.

Ingestion Task step properties

The Ingestion Task step properties display details about the current run of the Ingestion task of a taskflow.

When you click an Ingestion Task step, the **Properties** panel displays the following information:

Property	Description
Current Run	Displays the following details: - Status. Status of the mass ingestion task. - Start Date. Date and time when the mass ingestion task started. - End Date. Date and time when the mass ingestion task ended or was suspended. - Duration. Amount of time in seconds that the mass ingestion task ran before it completed or was suspended. - Started By. Name of the taskflow that called the mass ingestion task.
General	Displays the name of the Ingestion Task step.
Ingestion Task	Displays the name and type of the ingestion task that the Ingestion Task step contains. The type is set to **MI_TASK** for mass ingestion tasks.
Input Fields	Displays the following input fields: - Wait for Task to Complete. Indicates whether the Ingestion Task step waits for the mass ingestion task to complete. This value is always set to **true**. - Max Wait. Maximum time in seconds that the Ingestion Task step waits for the mass ingestion task to complete. This value is always set to **604800** seconds, which is 7 days. - Task Name. Name of the mass ingestion task. - GUID. Unique identification number for the mass ingestion task. - Task Type. Indicates the type of task that the Ingestion Task step contains. This value is always set to **MI_TASK** for mass ingestion tasks. - serviceProviderId. Service provider ID for the mass ingestion task. This value is always set to **filelistener**. - serviceProviderContextPath. Service provider context path for the mass ingestion task. This value is always set to **mitask/**.

Output Fields	Displays the following output fields if the Ingestion Task step completed successfully:
	- endTime. Date and time when the mass ingestion task completed.
	- failedFiles. Number of failed files in the mass ingestion task run.
	- jobNumber. ID for the mass ingestion task run.
	- messageText. Indicates whether the job completed normally. This value is always set to **Job completed normally**.
	- startTime. Date and time when the mass ingestion task started.
	- status. Indicates whether the Ingestion Task step completed completed. This value is always set to **Success**.
	- successFiles. Number of files that were successfully processed in the mass ingestion task.
	- fileDetails. Provides a list of objects with the following details:
	  - Duration. Amount of time in seconds that the mass ingestion task ran before it completed.
	  - Last modified. Date and time when the file was last modified.
	  - Path. Name of the file and the path where the file exists.
	  - Size. Size of the file in KB.
	  - Status. Indicates the status of the file transfer. This value is always set to **Success**.
	  - Transfer direction. Indicates the transfer direction for the file. You see one of the following values:
	    - DOWNLOAD. The file was downloaded from the source.
	    - UPLOAD. The file was uploaded to the target.

Fault	If the ingestion task faulted, you see the following properties: - Code. Indicates whether the fault is an error or a warning. - Reason. Cause of the error or fault.

Consider the following rules and guidelines when you monitor the Ingestion Task step:

- The **All Jobs** page does not display the ingestion task as a subtask. It always displays the number of rows processed as **0** if the taskflow includes an Ingestion Task step.
- If the ingestion task faulted and the **On Error** property was set to **Suspend Taskflow**, no information is displayed in the **Properties** panel.

Monitoring linear taskflows

To view detailed information about a linear taskflow instance, click the job name on the **All Jobs**, **Running Jobs**, or **My Jobs** page.

The following image shows the job details for a linear taskflow:
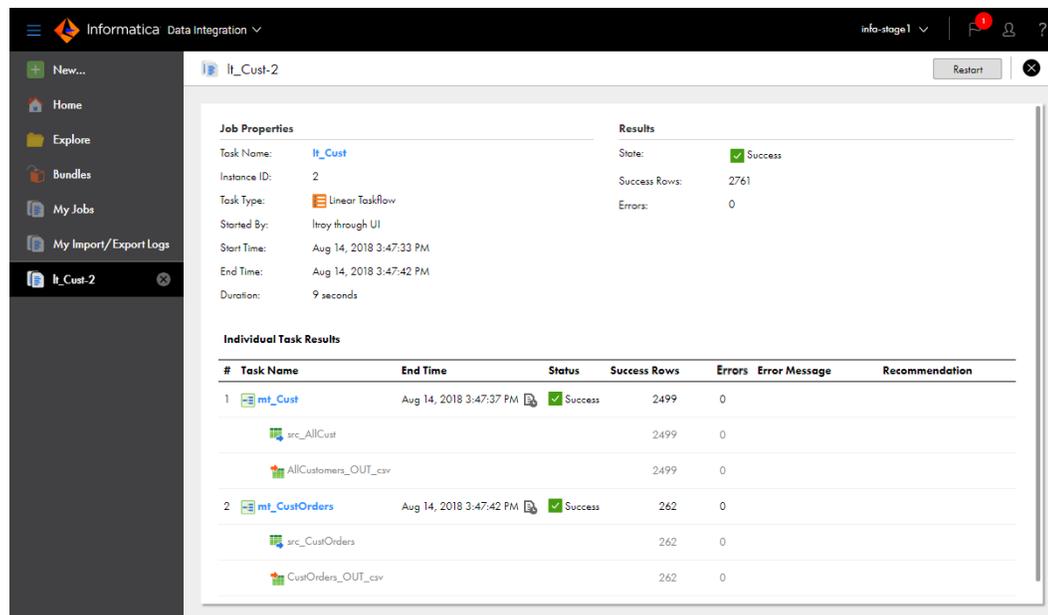
You view different details in each panel.

Job properties

The job properties for each linear taskflow instance display general properties about the instance.

The job properties include the following properties:

Property	Description
Task name	Name of the linear taskflow.
Instance ID	Instance number for the linear taskflow. For example, if you are looking at the third run of the taskflow, this field displays "3."
Task type	Task type, that is, Linear Taskflow.

Property	Description
Started by	Name of the user or schedule that started the job.
Start time	Date and time that the job was started.
End time	Date and time that the job completed or stopped.
Duration	Amount of time the job ran before it completed or was stopped.

Job results

The job results for each linear taskflow instance display the status of the job and success and error statistics.

The job results include the following properties:

Property	Description
Status	Job status. A job can have one of the following statuses: - Starting. The job is starting. - Running. The job is still running. - Success. The job completed successfully. - Warning. The job completed with errors. - Failed. The job did not complete because it encountered errors.
Success rows	Number of rows successfully written to the target.
Errors	Number of rows that were not written to the target.

Individual task results

The individual task results display results for individual subtasks in the linear taskflow instance.

For each subtask, the job details include the following properties:

Property	Description
Task name	Name of the subtask.
End time	Date and time that the subtask completed or stopped.
Status	Status of the subtask. A subtask can have one of the following statuses: - Queued. The job is queued on a Secure Agent, but it has not started yet. - Starting. The job is starting. - Running. The job is still running. - Success. The job completed successfully. - Stopped. The parent job has stopped running, so the subtask cannot start. - Warning. The job completed with errors. - Failed. The job did not complete because it encountered errors. - Aborted. The job was aborted. Applies to mass ingestion subtasks.

Property	Description
Success rows	Number of rows successfully written to the target. For each source and target in the subtask, this field displays the number of rows successfully read from the source and the number of rows successfully written to the target.
Errors	Total number of source error rows, target error rows, and transformation errors. For each source and target in the subtask, this field displays the number of rows that were not read from the source and the number of rows that were not written to the target.
Error message	Error encountered when writing to the target, if any.
Recommendation	Recommended action for errors.

Stopping and restarting mapping and task instances

You can stop and restart mapping and task instances on the **All Jobs**, **Running Jobs**, or **My Jobs** page. Restart a mapping or task instance on the **All Jobs** or **My Jobs** page or from the job details.

You can perform the following actions:

Stop a mapping or task instance.

You can stop a mapping or task instance that is starting, running, or queued.

To stop a mapping or task instance, open the **All Jobs**, **Running Jobs**, or **My Jobs** page. Then, click the **Stop** icon in the row that contains the job that you want to stop.

Restart a mapping or task instance.

You can restart a mapping or task instance that has run successfully or failed.

To restart a mapping or task instance on the **All Jobs** or **My Jobs** page, click the **Restart** icon in the row that contains the job that you want to restart.

To restart a mapping or task instance from the job details, click **Restart** at the top of the page.

Stopping, suspending, resuming, and restarting taskflow instances

You can stop, suspend, resume, and restart taskflow instances on the **All Jobs**, **Running Jobs**, or **My Jobs** page.

You can perform the following actions:

Stop a taskflow instance.

You can stop a taskflow instance that is in the suspended or running state.

To stop a taskflow instance, open the **All Jobs**, **Running Jobs**, or **My Jobs** page, and then click the **Stop** icon in the row that contains the taskflow instance.

Suspend a taskflow instance.

To suspend a taskflow instance, open the **All Jobs**, **Running Jobs**, or **My Jobs** page, and then click the **Suspend** icon in the row that contains the taskflow instance.

Resume a taskflow instance.

To resume a taskflow instance, open the **All Jobs**, **Running Jobs**, or **My Jobs** page, and then click the **Resume** icon in the row that contains the taskflow instance.

The taskflow resumes from the step at which it was suspended. If you resume a taskflow that was suspended because of a Data Integration task fault, the task re-runs.

For information about suspending a taskflow on a data task fault, see *Taskflows* in the Data Integration service help.

Restart a taskflow instance.

Users who have the Admin or Operator role can restart a taskflow from Monitor and Data Integration.

You can restart a taskflow that meets the following criteria:

- The taskflow was run at least once.
- The taskflow is in one of the following states:
  - Success
  - Failed
  - Suspended

You can also restart a taskflow instance that has the **Binding** property set to **Event**.

To restart a taskflow instance, open the **All Jobs**, **Running Jobs**, or **My Jobs** page, and then click the **Restart** icon in the row that contains the taskflow instance.

CHAPTER 3

Data Integration job log files

Data Integration generates log files to help you monitor running, failed, and completed jobs. You can access the log files from the **All Jobs**, **Running Jobs**, and **My Jobs** pages. You can also access the log files from the job details.

Data Integration generates the following types of log files:

Error rows file

Data Integration generates error rows files for synchronization task and masking task instances. An error rows file shows the rows that failed and the reason why each row failed. The error rows file includes the first 50 fields of a source error row.

For example, the following error appears in the error rows file when the task tries to insert two records with the same external ID into a Salesforce target:

```
Error loading into target [HouseholdProduct__c] : Error received from salesforce.com.
Fields [ExternalId__c]. Status code [DUPLICATE_VALUE]. Message [Duplicate external id
specified: 1.0].
```

Session log file

Data Integration generates a session log file for each job. If a job fails, download the log file to help you troubleshoot the job.

Execution plan

An execution plan shows the runtime Scala code that the Serverless Spark engine uses to run the data logic in an elastic mapping. You can use the Scala code to debug issues in the mapping.

Agent job log

An agent job log shows the logic that the Secure Agent uses to push the Spark execution workflow for an elastic job to an elastic cluster for processing.

Spark driver and Spark executor logs

Spark driver and Spark executor logs show the logic that the Serverless Spark engine uses to run an elastic job.

Initialization script log

If an initialization script runs on the cluster, the init script log shows the script output.

Cloud-init log

If an initialization script runs on the cluster, the cloud-init log contains information about how cluster nodes were initialized and bootstrapped. You can use the cloud-init log to check if any init scripts failed to run.

Note: You can view the cloud-init log only in an AWS environment.

Viewing an error rows file

You can view an error rows file using one of the following methods:

Preview error rows file

Preview the error rows file from the **All Jobs**, **Running Jobs**, or **My Jobs** page or from the job details. When you preview the error rows file, Data Integration generates a CSV file that contains the first 25 error rows and first 50 columns.

To preview an error rows file, your user role must have the "Access CDI Error Logs" feature privilege for Data Integration.

In Monitor, select **All Jobs** or **Running Jobs**, or in Data Integration, select **My Jobs**.

Perform either of the following steps:

- Click the **Preview Error Rows File** icon in the row that contains the task instance.
- Click the job name to open the job details, and then click **Preview Error Rows File** in the Results area.

View error rows file

To view the entire error rows file for a job, open the file in a text editor. If the error rows file contains Unicode data from the source, open the error rows file in an application that can display Unicode characters.

Error rows files are written to the following directory:

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

Error rows files are named with a timestamp that identifies the start time of the job. To identify the correct error rows file, use the start time of the job.

Downloading a session log file

You can download the log file for a job that is running, completed, or failed. Download the log file from the **All Jobs**, **Running Jobs**, or **My Jobs** page or from the job details.

1. In Monitor, select **All Jobs** or **Running Jobs**, or in Data Integration, select **My Jobs**.
2. Perform either of the following steps:
   - Click the **Download Log** icon in the row that contains the job.
   - Click the job name to open the job details, and then click **Download Session Log** in the Results area.

Viewing log files for elastic jobs

When you monitor an elastic job, you can view the execution plan, the session log, the agent job log, and the Spark driver and Spark executor logs. If an initialization script runs on the cluster, you can also view the init script logs and cloud-init logs.

To access a specific log file for an elastic job, you might have to download the log from Monitor or from your cloud platform.

The following table indicates whether you can download each log from Monitor, your cloud platform, or both:

Log	Download location
Execution plan	Monitor
Session log	Monitor
Agent job log	Monitor, cloud platform[1]
Spark driver log	Monitor, cloud platform[1]
Spark executor log	Monitor[2], cloud platform[1]
Initialization script log	Cloud platform
Cloud-init log	Cloud platform

1 *If you use a serverless runtime environment, downloading from your cloud platform is not available.*

2 *If you do not use a serverless runtime environment, downloading from Monitor is not available. Use the displayed log location to download the log from your cloud platform.*

Downloading the execution plan and session log

You can download the execution plan and the session log from the job results in Monitor.

Downloading advanced logs

You can download the agent job log and the Spark driver log using the **Actions** menu in the Spark task details in Monitor. You can also use the **Actions** menu to copy the advanced log location and navigate to the location on your cloud platform. You must navigate to the advanced log location to download Spark executor logs.

If you use a serverless runtime environment, use the **Actions** menu to download the advanced logs.

Downloading init script and cloud-init logs

To download an initialization script log or a cloud-init log, navigate to the following location on your cloud platform:

```
<Log location>/ClusterLogs/<cluster ID>/<timestamp>/<cluster node instance ID>/
```

The log location is the location for log files that is specified in the elastic configuration.

Note: You can download cloud-init logs only in an AWS environment.

Rules and guidelines

Refer to the following rules and guidelines when you view log files:

- If you download the Spark driver and agent job logs while the job is running, the files reflect the current state of the job.
- If you view a log file for an elastic mapping with a mapping name that begins with a number or a space, the mapping name is prefixed with an underscore in the log file.

  For example, when a log file references a mapping named `123Mapping`, the log files use the name `_123Mapping`.

CHAPTER 4

Monitoring Mass Ingestion jobs

You can monitor the progress, performance, and status of ingestion jobs from the Mass Ingestion and Monitor 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 created and 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 the Monitor service, monitor *all* types of ingestion jobs that any member of your organization created and deployed. You can view the following types of information:
  - Summary counts of ingestion jobs by task type and job state.
  - Recent jobs that require your attention because they have errors or warnings.
  - A list of all ingestion jobs by type, including the general job properties.

- From either the list of your jobs or list of all jobs, you can drill down to details for a specific job by clicking the job name. You can view additional job overview information, alert messages, performance statistics, and information about past job runs.

Note: A streaming ingestion job name matches the ingestion task name.

Monitoring your ingestion jobs

On the **My Jobs** page in Mass Ingestion, you can monitor the ingestion jobs for the tasks that you created and deployed.

The following image shows the **My Jobs** page with database ingestion jobs:

The following image shows the **My Jobs** page with streaming ingestion jobs:

For descriptions of the columns, see <u>"Job properties" on page 49</u>. These columns are the same as those shown for all ingestion jobs on the **All Jobs** tab of the **Mass Ingestion** page in Monitor service.

From the Action menu in a job row, you can perform some actions on the job.

Tip: 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.

- To find a job based on the job name, enter the job name, or any part of the name, in the *Find* text box. With a partial name, you do not need to include a wildcard. Do not include the following symbols: question mark (?), number sign (#), or ampersand (&). If you include one or more these symbols, no Find results are returned.

- To filter the list of jobs, click the **Filter** icon. Then click **Add Field** and enter filter criteria for one or more of the listed fields.

Monitoring all ingestion jobs

On the **Mass Ingestion** page in Monitor, you can monitor all ingestion jobs that were defined and deployed from the Mass Ingestion service, including database 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. Each button shows the total number of jobs with that state. It also lists the jobs that match the selected filters.

- 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 top of the **Overview** tab provides buttons for filtering the list of jobs by job type and state. Each button shows the number of jobs with that type or state.

The following example shows the **Overview** tab for database ingestion jobs with all status buttons displayed:
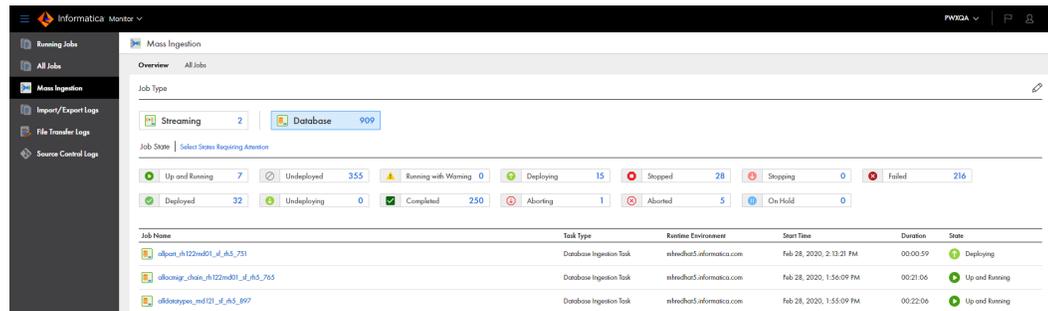
The following example shows the **Overview** tab for streaming ingestion jobs with all status buttons displayed:
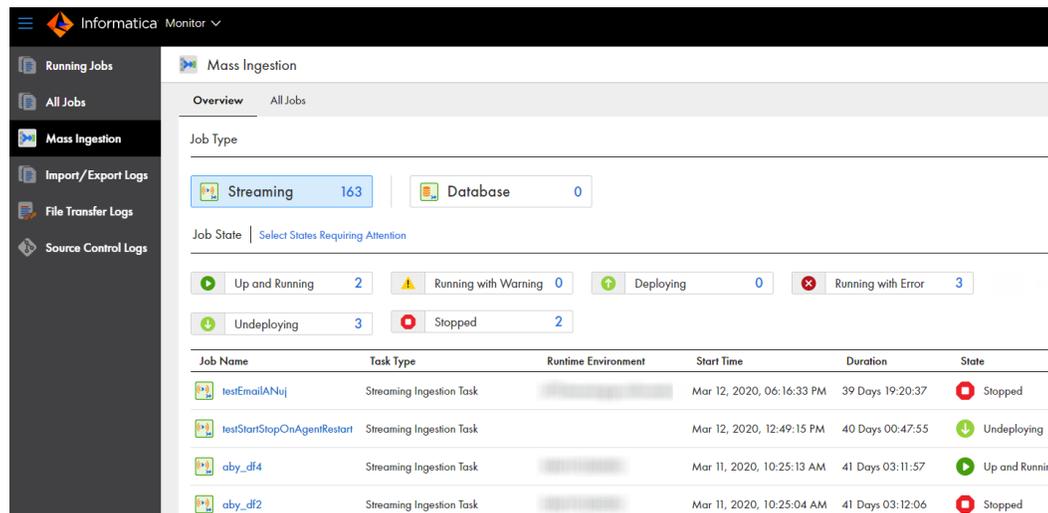
To control the state buttons that appear on the **Overview** tab, click the Edit (pencil) icon. Then in the **Reorder Job State** dialog box, select the **Visibility** check box next to each job state that you want to display buttons and jobs for.

To rearrange the order of the job state buttons, click the Edit icon. Then in the **Reorder Job State** dialog box, select a job state row and click the up or down arrow icon to move the row.

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

- To see only database ingestion jobs or only 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.

- To see the jobs that match a particular job state, click a state button. You cannot select multiple state buttons at the same time on the **Overview** tab. To clear the filter, click the selected state button again.

- To see the jobs with a state that might require your attention, click **Select States Requiring Attention** above the job state buttons. This option is displayed only after you select a job type. This option lists database ingestion jobs with the **Failed** or **Running with Warning** state, or lists streaming ingestion jobs with the **Running with Error** or **Running with Warning** state. To clear the filter, click **Select States Requiring Attention** again.

Note: All filters that you set on the **Overview** tab or in the **Reorder Job States** 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 of any type and in any state.

The following example shows the **All Jobs** tab for database ingestion jobs:
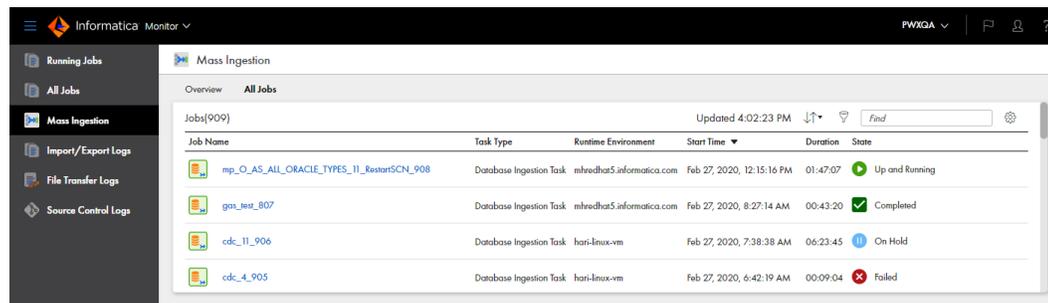
The following example shows the **All Jobs** tab for streaming ingestion jobs:
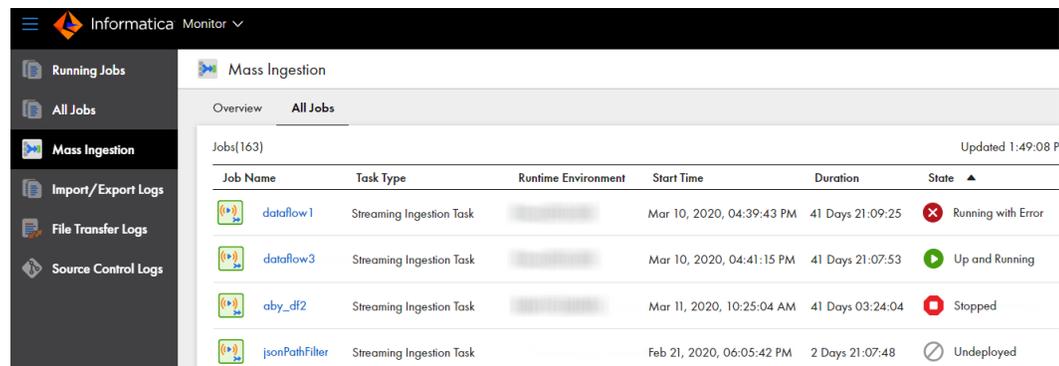
From the Actions menu at the right end a job row, you can perform some actions on the job, depending on the job state.

Tip: 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** icon and select a field to sort by.

- To find a job based on the job name, enter the job name, or any part of the name, in the *Find* text box. With a partial name, you do not need to include a wildcard. Do not include the following symbols: question mark (?), number sign (#), or ampersand (&). If you include one or more these symbols, no Find results are returned.

- To filter the list of jobs, click the **Filter** icon. Then click **Add Field** and enter filter criteria for one or more of the listed fields.

Job properties

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

The following table describes the job properties:

Property	Description
Job Name	The generated name of the ingestion job. **Note:** If you edit the name of the associated ingestion task, the job name remains the same.
Task Type	The type of ingestion task. This value must be either **Database Ingestion Task** or **Streaming Ingestion Task**.
Runtime Environment	The name of the runtime environment in which the job runs.
Start Time	For database ingestion jobs, the date and time when the job was deployed. For streaming ingestion jobs, the date and time when the job was deployed.
Duration	For 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 streaming ingestion jobs, the amount of time that the job has been running.
State	The current state of the job, such as Deploying, Up and Running, or Undeployed. The set of valid states vary by type of ingestion task. For more information, see *Overview tab* in "Streaming ingestion job details" on page 54 or the Job Overview tab in "Database ingestion job details" on page 49.

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 Monitor, 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.

Database ingestion job details

For database ingestion jobs, you can view job-specific details on the **Job Overview**, **Object Detail**, and **Alerts** tabs. Access these tabs by drilling down on a job from the **My Jobs** page in Mass Ingestion or from the **All Jobs** tab on the Mass Ingestion page in Monitor.

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 state, 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 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 state. 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 *task name-job instance number*.
Task Name	The name of the associated ingestion task.
Task Type	The type of task, which is Database Ingestion Task.
Load Type	The type of load operation that the job performs. Options are: - **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.
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.

Property	Description
State	The state of the job, 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 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	The number of records that were read from the source.
Records Written	The number of records that were successfully propagated to an Amazon S3, Apache Kafka, flat file, or Microsoft Azure Data Lake Storage target. For a Microsoft Azure SQL Data Warehouse target, the number of records written to the intermediate Microsoft Azure Data Lake Storage files. For a Snowflake target, the number of records written to the internal staging area that is created when the job runs.
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.
Log	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: - **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. To download a log to your local system, click the Download icon. **Note:** For initial load jobs, you can get the job log for a specific source object from the **Object Detail** tab.

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. For an initial load, you can also download the complete job log or the error log for a source object.

The following image shows the **Object Detail** tab for a database ingestion job:
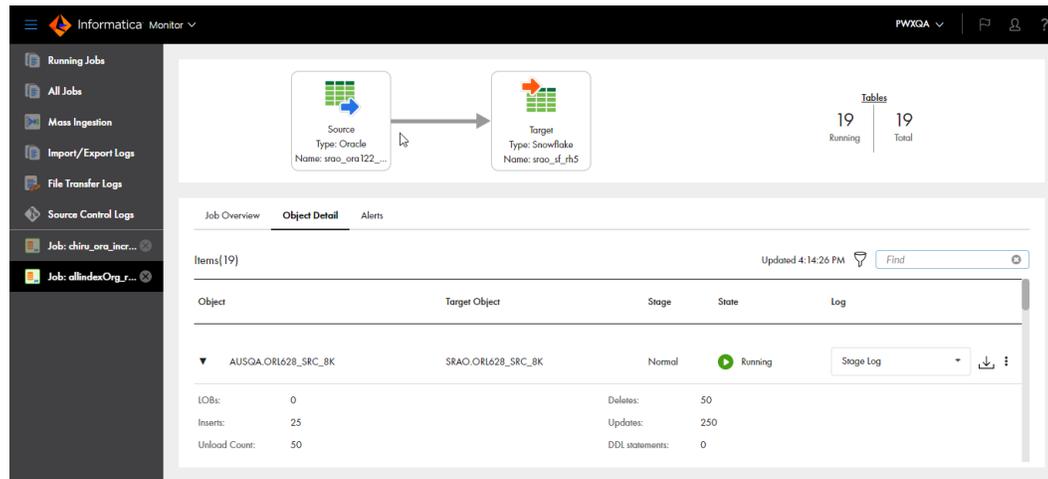
The following table describes the properties that are displayed for each object:

Column	Description
Object	The name of the source table for which data was propagated to the target.
	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: - 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	For an initial load job, the number of records that were successfully propagated 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. For a Microsoft Azure SQL Data Warehouse target, the number of records written to the intermediate Microsoft Azure Data Lake Storage files. For a Snowflake target, the number of records written to the internal staging area that is created when the job runs.

Column	Description
Task Duration	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. When a job runs, it uses a separate subtask to process each source table.
Stage	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. The stage can be one of the following values: - **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.
State	The state of the job subtask for the source object. **Note:** If the job stops running, the subtask state reflects the state last collected before the job ended. For example, the job might be aborted but the subtask is in a Running state. The state can be one of the following values: - **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. **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.

Column	Description
Log	For initial load jobs, you can download the job execution log for a source object. Select one of the following log types: - **Complete Log**. The complete log from job execution. - **Error**. The log that contains error messages. This log type is available only for a Failed job. 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. To download a log locally, click the Download icon. **Notes:** - For incremental load jobs, you must get the complete log for the entire job run from the **Job Overview** tab. - If you undeployed the job, you can download the log for a table only if the associated task has not been deleted.
Actions menu > Resync	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.

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.

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.

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 the **Monitor** service.

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 Mass Ingestion Streaming task.
Task Location	The project or project folder that contains the Mass Ingestion Streaming 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: - 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.

Property	Description
Start Time	The date and time when the job was deployed.
Runtime Environment	Name of the runtime environment that the job uses to run.
Download Log	Level of log that you want to download for a running job. You can download one of the following logs: - Complete Log. 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. To download a log to your local system, click the **Download** icon.

Alert tab

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

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

The following table describes the job alert properties:

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

Performance tab

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

The following image shows the **Performance** tab for a streaming ingestion job:
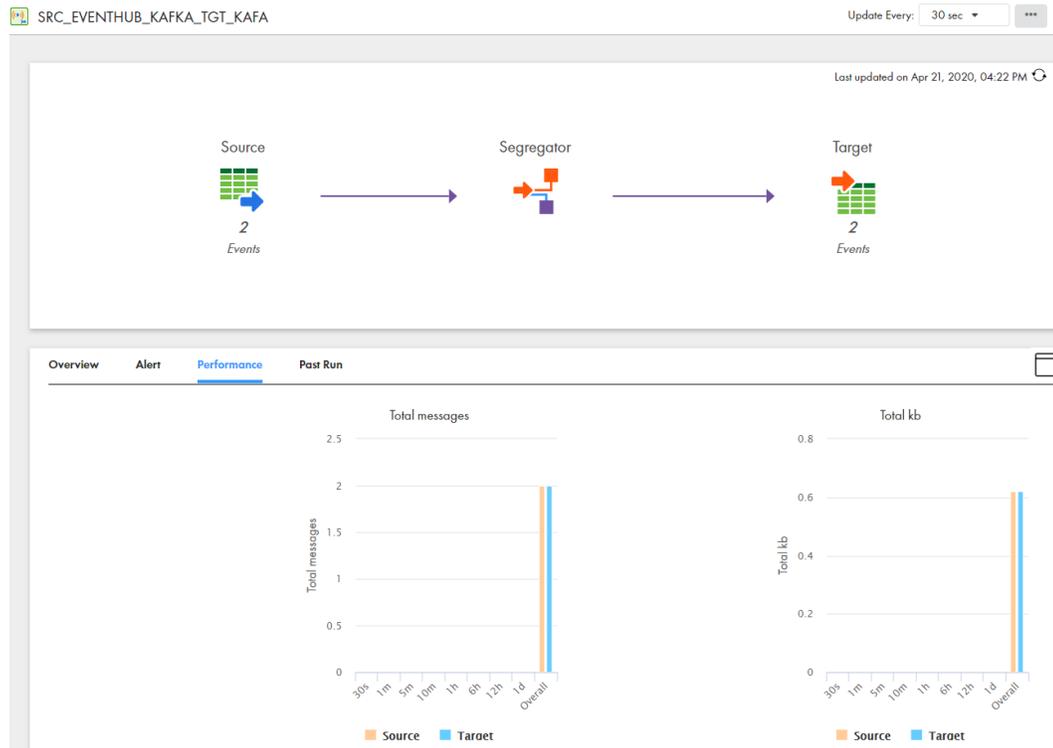
The following table describes the job performance properties:

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.

CHAPTER 5

Monitoring Data Accelerator for Azure jobs

You can monitor the jobs that you are running or have run. A job is an instance of a mapping, task, or taskflow. Each time that you start an Azure data sync task, Data Accelerator for Azure creates a job to run it.

You can monitor Azure data sync jobs on the following pages:

All Jobs page

Lists all jobs that are running or that have run in your organization. You can stop and restart jobs on this page. If you open the subtask for an Azure data sync job, you can download the log file.

To open the **All Jobs** page, in Monitor, select **All Jobs**.

Running Jobs page

Lists all jobs that are running or have completed within the last five minutes. You can stop and restart jobs on this page. If you open the subtask for an Azure data sync job, you can download the log file.

To open the **Running Jobs** page, in Monitor, select **Running Jobs**.

My Jobs page

Lists all jobs that were started by the currently logged in user. You can stop and restart jobs on this page. If you open the subtask for an Azure data sync job, you can download the log file.

To open the **My Jobs** page, in Data Accelerator for Azure, select **My Jobs**.

Job details page

Displays detailed information about a specific job. You can view the details for a job that has completed or failed. When you view details for an Azure data sync job, you can restart the job and download the log file.

To open the job details for an Azure data sync job, open the subtask for the job on the **All Jobs**, **Running Jobs**, or **My Jobs** page, and then click the job name.

Job properties

The **All Jobs**, **Running Jobs**, and **My Jobs** pages display the job properties such as the name, start time, and status. You can right-click the column heading area to display or hide specific properties.

Each Azure data sync job is named `<task name>-<instance number>`. For example, the first time you run Azure data sync task ads_LoadOrders, the job is named ads_LoadOrders-1. The second time you run the task, the job is named ads_LoadOrders-2.

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

Property	Description
Instance name	Name of the job in the following format: `<task name>-<instance number>`
Location	Project and folder path where the task exists.
Subtasks	The number of subtasks. Each Azure data sync job contains one subtask, which is a mapping. To open the subtask view, click the entry in this column. To close the subtask view, click **Data Integration** from the menu at the top of the page.
Start time	Date and time that the job was started.
End time	Date and time that the job completed or stopped. Does not apply to running jobs.
Rows processed	Total number of rows that the job has currently processed. This value includes the number of rows successfully written to the target as well as the number of error rows for the source and target.
Status	Job status. An Azure data sync job can have one of the following statuses: - Starting. The job is starting. - Running. The job is still running. - Success. The job completed successfully. - Warning. The job completed with errors. - Failed. The job did not complete because it encountered errors. **Tip:** If the job status is Failed, you can hover the cursor over the job status to view and copy the error message.

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

Property	Description
Asset name	Name of the asset that is associated with the job. For Azure data sync jobs, this column displays the name of the associated Azure data sync task. In the subtask view, this column displays the name of the underlying mapping.
Asset type	Type of asset that is associated with the job. For Azure data sync task jobs, this column displays "Azure Data Sync Task." In the subtask view, this column displays "Mapping."
Runtime environment	Runtime environment in which the job ran.
Duration	Amount of time the job ran before it completed or was stopped.
Success rows	Number of rows successfully written to the target.

Property	Description
Failure rows	Number of rows that were not written to the target.
Started by	Name of the user or schedule that started the job.
Error message	Error message, if any, that is associated with the job.

Viewing details for an Azure data sync job

To view detailed information about a specific Azure data sync job, open the **All Jobs**, **Running Jobs**, or **My Jobs** page and click the number of subtasks in the row that contains the job. Then click the instance name in the subtasks view.

The job details displays the following properties for Azure data sync jobs:

Property	Description
Task name	Name of the mapping that the Azure data sync task is based on.
Instance ID	Number of times the mapping has been run.
Task type	Task type. For Azure data sync jobs, this field displays "Mapping" because each Azure data sync job is based on an underlying mapping.
Started by	Name of the user or schedule that started the job.
Start time	Date and time that the job was started.
End time	Date and time that the job completed or stopped.
Duration	Amount of time the job ran before it completed or was stopped.
Runtime environment	Runtime environment in which the job ran.
Secure Agent	Secure Agent that ran the job.
Context name	Name of the Azure data sync task that is associated with the job.
Status	Job status. An Azure data sync job can have one of the following statuses: - Starting. The job is starting. - Running. The job is still running. - Success. The job completed successfully. - Warning. The job completed with errors. - Failed. The job did not complete because it encountered errors.
Success rows	Number of rows successfully written to the target.
Errors	Number of rows that were not written to the target.

Property	Description
Session log	Allows you to download the session log file. By default, Data Accelerator for Azure 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`
Error message	Error message that is associated with the job if the job failed.
Individual source/ target results	The name, number of success rows, and number of errors for the source and target in the Azure data sync task. If the source or target has errors, this area also displays the error messages and recommended actions.
Job payload	Displays the runtime parameters that were used to run the underlying mapping and associated mapping task. Used for troubleshooting the job.

Stopping and restarting Azure Data Sync jobs

You can stop and restart Azure data sync jobs. Stop a job on the **All Jobs**, **Running Jobs**, or **My Jobs** page. Restart a job on the **All Jobs** or **My Jobs** page or from the job details.

You can perform the following actions:

Stop a job.

You can stop an Azure data sync job that is starting, running, or queued.

To stop a job, open the **All Jobs**, **Running Jobs**, or **My Jobs** page. Then click the **Stop** icon in the row that contains the job that you want to stop.

Restart a job.

You can restart an Azure data sync job that has run successfully or failed.

To restart a job on the **All Jobs** or **My Jobs** page, click the **Restart** icon in the row that contains the job that you want to restart.

To restart a job from the job details, click **Restart** at the top of the page.

Downloading a log file

You can download the log file for an Azure data sync job that has completed or failed. Download the log file from the subtask view of the **All Jobs**, **Running Jobs**, or **My Jobs** page or from the job details.

1. In Monitor, select **All Jobs** or **Running Jobs**, or in Data Accelerator for Azure, select **My Jobs**.
2. Click the number of subtasks in the Subtasks column to open the subtask view.
3. Perform either of the following steps:

- Click the **Download Log** icon in the row that contains the job.

- Click the job name to open the job details, and then click **Download Session Log** in the Results area.

CHAPTER 6

Monitoring Data Profiling jobs

You can monitor the jobs that are running or have run. A job is an instance of a task. Each time that you run a data profiling task, Data Profiling creates a job to run it.

You can monitor data profiling jobs on the following pages:

All Jobs page

Lists all jobs that are running or that have run in your organization. You can stop and restart jobs on this page. If you open the subtasks for a data profiling job, you can view the subtask details.

To open the **All Jobs** page, in Monitor, select **All Jobs**.

Running Jobs page

Lists all jobs that are running or have completed within the last five minutes. You can stop and restart jobs on this page. If you open the subtasks for a data profiling job, you can view the subtask details.

To open the **Running Jobs** page, in Monitor, select **Running Jobs**.

My Jobs page

Lists all jobs that were started by the currently logged in user. You can stop and restart jobs on this page. If you open the subtasks for a data profiling job, you can view the subtask details.

To open the **My Jobs** page, in Data Profiling, select **My Jobs**.

Job properties

The **All Jobs**, **Running Jobs**, and **My Jobs** pages display the job properties such as the name, start time, and status. You can right-click the column heading area to display or hide specific properties.

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

Property	Description
Instance name	Name of the job.
Location	Project and folder path where the task exists.
Subtasks	Shows the number of subtasks created for the job. To view the job properties, click the entry in this column. To close the subtask view, click Data Integration from the menu at the top of the page.

Property	Description
Start time	Date and time that the job was started.
End time	Date and time that the job completed or stopped. Does not apply to running jobs.
Rows processed	Number of source rows processed by the job.
Status	Job status. A job can have one of the following states: - Starting. The job is starting. - Running. The job is still running. - Success. The job completed successfully. - Warning. The job completed with errors. - Failed. The job did not complete because it encountered errors. When a job is in Running status, a Stop icon appears in the Status property. Click the icon to stop the job.

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

Property	Description
Asset name	Name of the data profiling task.
Asset type	Type of asset that is associated with the job. For data profiling task jobs, this column displays "Data Profiling Task."
Runtime environment	Runtime environment in which the job ran.
Duration	Amount of time the job ran before it completed or was stopped.
Success rows	Number of source rows processed successfully by the job.
Failure rows	Number of source rows that failed for the job.
Started by	Name of the user or schedule that started the job.
Error message	Error message, if any, that is associated with the job.

When you click an instance name, the job details for the data profiling job appears. The job details include the instance name, asset name, asset type, started by, start time, end time, duration, state, and error message.

Subtask types

When you run a data profiling task, a job instance is created for the profile run. Data Profiling creates subtasks based on the options that you choose for the profile run.

To view the subtasks for a job, click the subtask link in the **Subtasks** column. The subtasks created for the job appears.

Data Profiling creates the following subtasks based on the profile options that you choose for the profile run:

Subtask Name	Description
Fetching the source row count	Fetches the row count for the profile run. This subtask is created only once for a profile run.
Generating data profiling mappings	Gathers statistics based on the advanced options that you choose for the profile run. This subtask is created only once for a profile run.
s_profiling	Runs the profile and saves the results to a secure and encrypted staging area. This subtask is created multiple times based on the advanced options for the data profiling task. For information about the advanced options, see *Data Profiling*. The staging area is a temporary location in the Informatica Intelligent Cloud Services repository.
Loading data from staging area to metric store	Loads the data from the staging area to the metrics store. This subtask is created only once for a profile run. The metrics store is located in the Informatica Intelligent Cloud Services repository where the encrypted profile results are stored. When you choose a profile run number to view the results, the service gets the profile results from the repository.

To view the details for each subtask, click the subtask link in the **Instance Name** column. The subtask details include instance name, asset name, asset type, started by, start time, end time, duration, state, and error message. Additional details appear for the s_profiling subtask which include processed rows, success rows, error rows, and session log. Click the **Download Session Log** to download the session log file.

Stopping and restarting data profiling jobs

You can stop and restart data profiling jobs. Stop a job on the **All Jobs**, **Running Jobs**, or **My Jobs** page. Restart a job on the **All Jobs** or **My Jobs** page or from the job details.

You can perform the following actions:

Stop a job.

You can stop a data profiling job that is running or queued.

To stop a job, open the **All Jobs**, **Running Jobs**, or **My Jobs** page. Then click the **Stop** icon in the row that contains the job that you want to stop.

Restart a job.

You can restart a data profiling job that has failed.

To restart a job on the **All Jobs** or **My Jobs** page, click the **Restart** icon in the row that contains the job that you want to restart.

To restart a job from the job details, click **Restart** at the top of the page.

Downloading a log file

You can download the log file for a data profiling job that has completed or failed.

1. In Monitor, select **All Jobs** or **Running Jobs**, or in Data Profiling, select **My Jobs**.
2. Click the number of subtasks in the **Subtasks** column to open the subtask view.
3. Click **s_profiling** subtask.
4. In the **Results** area, click **Download Session Log**.

CHAPTER 7

Monitoring imports and exports

You can monitor the imports and exports that are running or have run in your organization. Each time that you import or export objects, Informatica Intelligent Cloud Services creates a log entry for the import or export.

You can view the import and export logs on the following pages:

Import/Export Logs page

Lists all imports or exports that are running or have run in your organization. If you view exports, you can download the export files from this page.

To open the **Import/Export Logs** page, open Monitor and select **Import/Export Logs**.

My Import/Export Logs page

Lists all imports or exports that were started by the currently logged in user. If you view your exports, you can download the export files from this page.

To open the **My Import/Export Logs** page, select **My Import/Export Logs**.

Details page

Displays detailed information about a specific import or export instance. You can also download the import or export instance log file. The log contains instance and object-level details. Logs are available for download for seven days. If you view details for an export instance, you can download the export file from this page.

To open the details page, click an import or export instance on the **Import/Export Logs** or **My Import/Export Logs** page.

Import and export log properties

To view properties for an import or export instance, click the **Import** or **Export** tab at the top of the **Import/Export Logs** or **My Import/Export Logs** page. By default, these pages list the import and export instances that were started within the last 7 days.

The following properties are displayed for each import or export instance:

Property	Description
Instance name	Name of the import or export instance.
Start time	Date and time that the import or export was started.

Property	Description
User name	Name of the user that started the import or export.
Start method	Method in which the import or export was started. An import or export instance can have one of the following start methods: - UI. The import or export was started from the user interface. - API. The import or export was started from the REST API.
Status	Status of the import or export, such as completed successfully or failed.

For successful export jobs, you can download the export file by clicking the **Download** icon in the row that contains the export instance.

Customizing import and export logs

You can decide which properties to display on the **Import/Export Logs** and **My Import/Export Logs** pages. You can sort and filter import or export instances. You can also drill down on an import or export instance to view detailed information.

The following image shows the properties that are displayed by default in Data Integration:
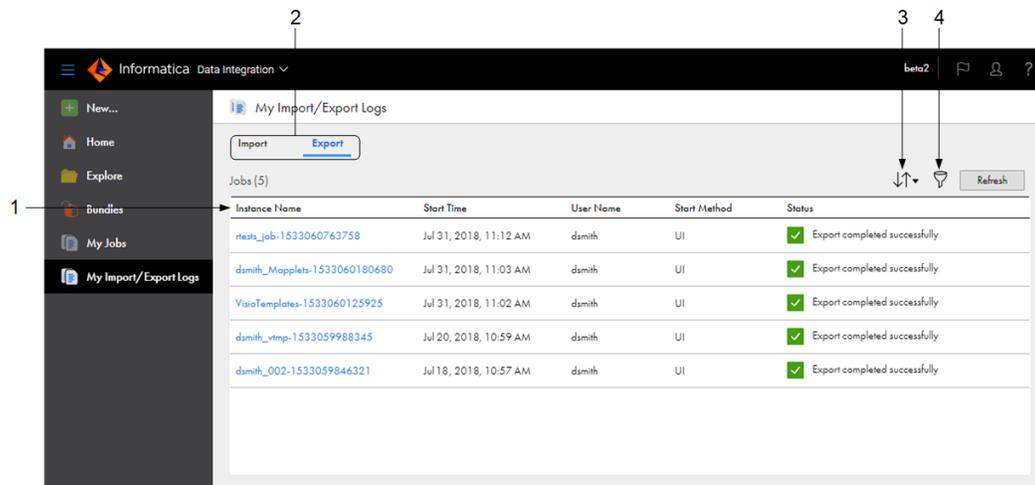
1. Column heading area. Right-click this area to add and remove columns.
2. Tabs to display either import logs or export logs
3. Sort icon
4. Filter icon

You can customize the import and export logs in the following ways:

Display, hide, or rearrange properties.

To display or hide specific properties, right-click the column heading area and check or uncheck the properties.

To rearrange the columns, click a column heading and drag it to a different location.

Sort import or export instances.

To sort the import or export instances, click the column heading for the property that you want to sort by. For example, to list the most recent jobs first, click the Start Time column. The arrow in the column

heading indicates the sort order, either ascending or descending. To reverse the sort order, click the column heading again.

You can also sort import or export instances by clicking the **Sort** icon and selecting the column name.

Filter import or export instances.

You can filter import or export instances using one or more filter conditions. To filter the instances that appear on the page, click the **Filter** icon.

To specify a filter condition, click **Add Field**, select the filter field, and then enter the value. For example, to view only successful exports or imports, select the Status column, and select "Success." You can select Instance Name or Status as the filter field.

To add another condition to the filter, click **Add Field** again. For example, to view failed instances named "SalesProject_1533326631921," select the following filter columns and values:

- Instance Name: SalesProject_1533326631921
- Status: Failed

To remove all applied filters, click the **Remove Filter** icon.

Viewing details for an import or export

To view detailed information about an import or export, click the **Import** or **Export** tab at the top of the **Import/Export Logs** or **My Import/Export Logs** page. Then click the import or export instance name.

The following image shows the details for an export instance:

The details include the following properties:

Property	Description
Name	Name of the import or export instance.
Type	Instance type, either Import or Export.
Start time	Date and time that the import or export was started.
End time	Date and time that the import or export completed or stopped.
Started by	Name of the user that started the import or export.
Start method	How the import or export was started. An import or export can have one of the following start methods: - UI. The import or export was started from the user interface. - API. The import or export was started from the REST API.
Status	Status of the import or export, such as completed successfully or failed.
Source organization	For imports, the name of the organization from which assets were imported.
Export file	For exports, name of the export file that was generated. Click the file name to download the file.
Import log / Export log	Detailed log file containing instance and object-level properties. Click **Download Import Log** or **Download Export Log** to download the file.
Imported assets / Exported assets	Name, type, location, description, and status of each asset that is imported or exported. The imported assets or exported assets list is available for one day.

Downloading an export file

When an export has completed, you can download the export file. Download the export file from the **Import/Export Logs** or **My Import/Export Logs** page or from the export instance details.

1. Select **Import/Export Logs** to open the import/export logs for the organization, or select **My Import/Export Logs** to open your import/export logs.
2. Click the **Export** tab.
3. Perform either of the following steps:
   - Click the **Download** icon in the row that contains the export instance.
   - Click the instance name to open the export instance details, and then click the export file name.

Downloading an import or export log

When an import or export instance is complete, you can download the log file. Download the log file from the import or export instance details.

1. Select **Import/Export Logs** to open the import/export logs for the organization, or select **My Import/Export Logs** to open your import/export logs.

2. Click the **Import** or **Export** tab.

3. Click the instance name to open the import or export instance details.

4. Click **Download Import Log** or **Download Export Log**.

CHAPTER 8

Monitoring file transfer jobs

You can monitor file transfer and file listener jobs that are running or have run in your organization.

To monitor file transfer jobs, click **File Transfer Logs**.

Monitoring AS2 file transfers

Whenever you send files through an AS2 connection or your AS2 server receives files, the file transfer service generates a log file. You can access the AS2 file transfer logs in Monitor.

You can view AS2 file transfer logs on the following pages:

File Transfer Logs page

> Lists all file transfers that are in progress or have completed.

Details page

> Displays detailed information about a specific file transfer log. You can also download the file transfer log file.

Viewing AS2 file transfer logs

To view a list of the AS2 file transfers that have completed, in Monitor, click **File Transfer Logs**.

To filter the list for AS2 file transfer logs, select **AS2 Logs**.

The **File Transfer Logs** page lists AS2 file transfer logs for send and receive file transfers and provides status information, as shown in the following image:

The **File Transfer Logs** page includes the following properties for each AS2 file transfer log:

Property	Description
Instance Name	Name of the AS2 file transfer instance.
Asset Type	Type of AS2 file transfer log. For sending files to remote AS2 servers, the type is AS2 Send. For receiving files from remote AS2 servers, the type is AS2 Receive.
Command	Type of activity, such as transfer, login, or MDN received.
End Time	Date and time that the file transfer ended.
# of Files	Number of files included in the transfer.
Sender/Receiver ID	For AS2 Send file transfers, the AS2 ID of the receiver. For AS2 Receive file transfers, the AS2 ID of the sender.
MDN	Status of the MDN message. For an AS2 Send, the status can be one of the following values: - Received. MDN is received. - Email. MDN delivery mode is through email. - URL. MDN delivery mode is through a URL. - None. MDN is not requested. For an AS2 Receive, the status can be one of the following values: - Pending. MDN has not been sent yet. - Sent. MDN is sent to the client. - None. MDN is not requested.
Status	Status of the AS2 file transfer, such as completed successfully or failed. To quickly find out why a transfer failed, rest your cursor over the Failed icon

Viewing details for an AS2 file transfer

To view detailed information about an AS2 file transfer including a list of transferred files, click the instance name on the **File Transfer Logs** page.

When you click the instance name, a detailed description of the file transfer opens. The type of information on the page depends on whether the instance type is AS2 Send or AS2 Receive and on the command type.

The log contains details including message properties and the status of each file included in the an AS2 file transfer. For an AS2 Send transfer, you can download the log file to help you troubleshoot the job.

To open the details page for an AS2 file transfer, click the instance name on the **File Transfer Logs** page. The following image shows the details for an AS2 Send transfer:
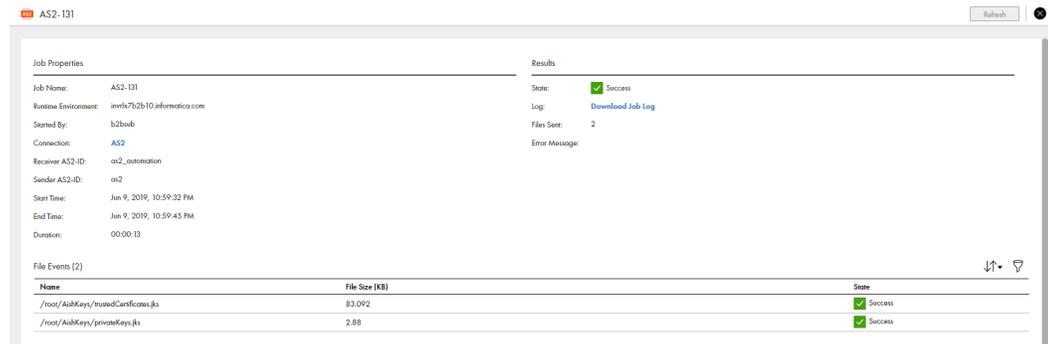
Details for AS2 Send file transfers

For an AS2 Send transfer, the log details include information about the job properties, job results, and files included in the job.

Job properties

Job details include the following properties:

Property	Description
Job Name	Name of the job.
Runtime Environment	Runtime environment that contains the Secure Agent used to run the job.
Started By	The user who initiated the job or name of the file listener that initiated the job.
Connection	AS2 connection used to send the files.
Receiver AS2-ID	ID of the remote AS2 server that received the files.
Sender AS2-ID	ID of the AS2 server that sent the files.
Start Time	Date and time that the file transfer started.
End Time	Date and time that the file transfer ended.
Duration	Amount of time that the file transfer ran.

Results

Job result information includes the following properties:

Property	Description
State	Status of the job, such as Success or Failed.
Log	Link to download the job log.

Property	Description
Files Sent	Number of files sent in the job.
Error Message	Error message if the job failed.

Files

Details about the files sent to the AS2 server include the following properties for each file:

Property	Description
Name	Full path and name of the file.
File Size	Size of the file, in kilobytes.
Status	Status of the file transfer, such as Success or Failed.

Details for AS2 Receive file transfers

For an AS2 Receive type of file transfer, the log details include information about the event properties, message properties, and files included in the file transfer.

Event details include the following properties:

Property	Description
Server IP	IP address of the AS2 server receiving the files.
Server Port	Port number for the AS2 server receiving the files.
Receiver AS2-ID	ID used by the recipient.
Remote IP	IP address of the client sending the files.
Remote Port	Port number for the client sending the files.
Sender AS2-ID	Name or ID used by the sender.
Username	User name of the file server user.
Command	Type of activity, such as transfer, login, or MDN received.
Subject	Subject of the message.
Encryption Algorithm	Algorithm used to encrypt the message.
Signature Algorithm	Algorithm used to sign the message.
Compressed	Indicates whether the message is compressed.

Property	Description
Content Type	Content type of the message.
Message ID	ID of the message.

MDN details include the following properties:

Property	Description
MDN Type	Whether the MDN is asynchronous or synchronous.
MDN Signed	Indicates whether the MDN is signed with a digital signature.
MDN Sent	Whether the MDN was sent.
MDN Delivery	Method used to deliver the MDN such as email, URL, job log, or file.
MDN	Link to view the MDN.
MIC	The message integrity check code.
MIC Algorithm	The MIC algorithm used for the signature.
State	Status of the transfer, such as Success or Failed.
Error Message	Error message if the transfer failed.

Details about the files include the following properties for each file:

Property	Description
Name	Full path and name of the file.
File Size	Size of the file, in kilobytes.
Status	Status of the file transfer, such as Success or Failed.

Monitoring SFTP file transfers

Whenever you exchange files with remote partners, the file transfer service generates a log file. You can access the SFTP file transfer logs in Monitor.

You can view SFTP file transfer logs on the following pages:

File Transfer Logs page

When you view all file transfer logs on the **File Transfer Logs** page, the page shows logs of all SFTP upload and download commands. When you filter the view to SFTP logs, the page shows logs of all SFTP commands, for example, login, connect, and disconnect.

Event Properties page

Displays detailed information about a specific file transfer log.

Viewing SFTP file transfer logs

To view a list of the SFTP file transfer logs, in Monitor, click **File Transfer Logs**.

To filter the list for SFTP file transfer logs, select **SFTP Logs**.

The **File Transfer Logs** page lists SFTP file transfer logs and provides status information. When the **File Transfer Logs** page is filtered for SFTP logs, it shows the SFTP command, as well.

The following image shows a **File Transfer Logs** page that is filtered for SFTP logs:

When the **File Transfer Logs** page is filtered for SFTP logs, it includes the following properties for each SFTP log:

Property	Description
Instance Name	Name of the SFTP file transfer instance.
Asset Type	SFTP Server.
Command	Command type, such as connect, login, or upload.
End Time	Date and time that the action executed by the command ended.
# of Files	Number of files included in the action.
Status	Status of the action, success or failed.

Viewing details of an SFTP file transfer log

To view detailed information about an SFTP file transfer log, click the instance name on the **File Transfer Logs** page.

When you click the instance name, a detailed description of the SFTP event opens. The information on the page depends on the SFTP command type.

The following image shows the details of an Upload command:
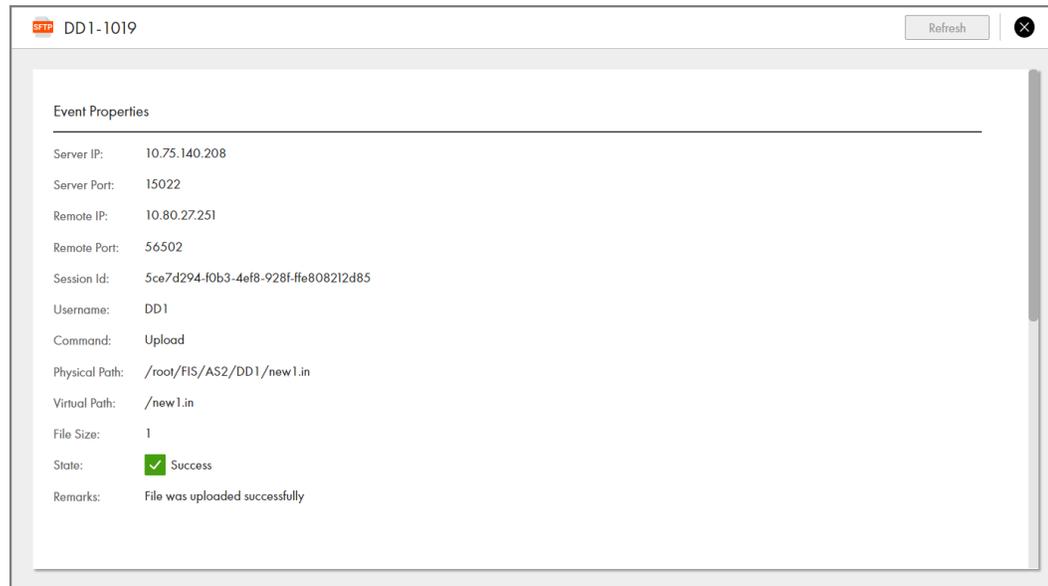
Details of an SFTP event

SFTP event properties depend on the SFTP command type. For example, physical path and virtual path don't apply to the Login command and are left blank in logs of Login commands.

Event details include the following properties:

Property	Description
Server IP	IP address of the SFTP server where the command was executed.
Server Port	Port number for the SFTP server where the command was executed.
Remote IP	IP address of the client where the command was initiated.
Remote Port	Port number for the client where the command was initiated.
Session ID	ID of the SFTP session.
Username	User name of the file server user.
Command	Type of SFTP command.
Physical Path	Absolute path to the file on which the command was executed in the runtime environment.

Property	Description
Virtual path	Relative path to the file on which the command was executed in the home directory of the file server user.
File size	Size of the file on which the command was executed.
State	Status of the action executed by the command.
Remarks	Additional details about the action executed by the command.

Monitoring HTTPS file transfers

When you exchange files with remote partners, the file transfer service generates a log file. You can access the HTTPS file transfer logs in Monitor.

View the HTTPS file transfer logs on the **File Transfer Logs** page. When you filter the page to view only HTTPS logs, the **File Transfer Logs** page displays the logs for all HTTPS commands, for example, login, logout, delete, rename, upload, and download.

Viewing HTTPS file transfer logs

To view the file transfer logs, in Monitor, click **File Transfer Logs**. To filter the list for HTTPS file transfer logs, select **HTTPS Logs**. The HTTPS file transfer logs shows the status of the file transfer and the HTTPS command.

The following image shows a sample **File Transfer Logs** page filtered for HTTPS logs:

The **File Transfer Logs** page includes the following properties for each HTTPS file transfer log:

| Property	Description
Instance Name	Name of the HTTPS file transfer instance.
Asset Type	HTTPS Server.
Command	Command type, such as login, upload, or rename.
End Time	Date and time when the action executed by the command ends.
# of Files	Number of files included in the action.
Status	Status of the action, such as Success or Failed.

Viewing details of HTTPS file transfer log

To view detailed information about an HTTPS file transfer log, click the instance name on the **File Transfer Logs** page. When you click the instance name, details about the HTTPS event are displayed. The information on the page varies based on the HTTPS command type.

The following image shows the details of a sample Upload command:
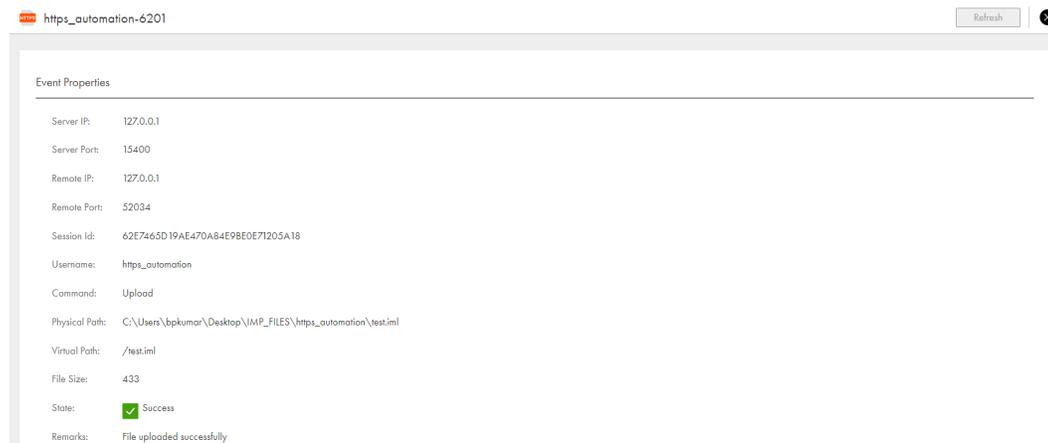
Monitoring file listeners

You can monitor executions of a file listener and the events that occur on each run job of the file listener. File listener log entries are listed on the file transfer logs page in Monitor.

1. In Monitor, select **File Transfer Logs**.

    The **File Transfer Logs** page lists file listener logs.

2. To filter the list for file listener logs, select **File Listener Logs** in the File Transfer Logs page.

3. Click a file listener.

    The file listener job details appears. This page includes details of the file listener job, results of the file listener monitoring, and details of file events that occurred in the defined location.
    The following image shows the state of the file listener and details of file events:
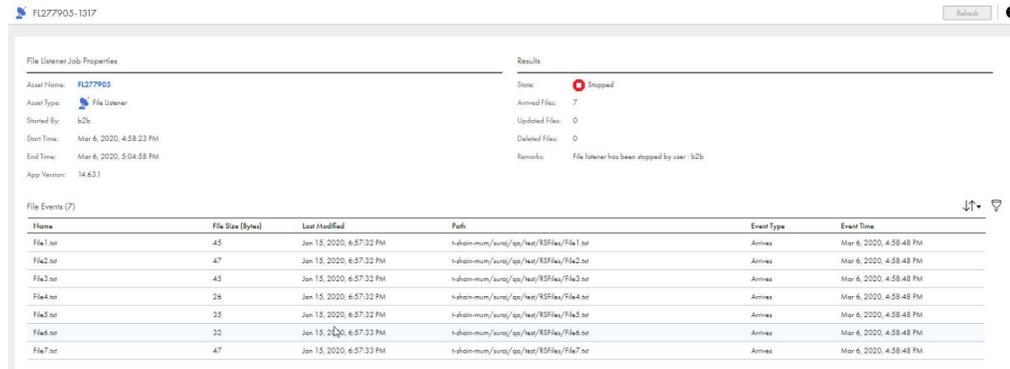
File listener status

The file listener creates a file listener job and displays in the File Transfer Logs page. The file listener listens to files in the specified folder and records events in the File Listener Job Properties page. The status field in the file listener job properties indicates the change in the file listener job.

The following table describes the file listener statuses:

Status	Reason
Failed	The file listener did not run successfully.
Completed	The file listener completed the run successfully.
Stopped	The stopped state occurs for the following reasons: - The file listener is stopped manually. - A scheduled run starts while the file listener is running after a manual start. A manually started job stops when a scheduled run begins. - The file listener stops automatically when the agent restarts on the same application version. The file listener also stops running on an application version if an updated application version is available.
Running	The file listener run is in progress.

Monitoring integration APIs

Integration Logs display the status of integration APIs such as Collect and Release commands used by external applications such as B2B Gateway to invoke file servers and transfer tasks. To filter the list of integration logs on the **File Transfer Logs** page, select **Integration Logs**. The integration logs are not displayed under the **All Logs** filter. You must use the **Integration Logs** filter to view the integration logs.

Viewing integration logs

To view integration logs in Monitor, click **File Transfer Logs**. To filter the list of integration logs, select **Integration Logs**.

The following image shows a sample **File Transfer Logs** page filtered for integration logs:

The **File Transfer Logs** page displays the following properties for each integration log:

Property	Description
Instance Name	Name of the integration log.
Asset Type	API or a file transfer task.
Command	Command type, such as Collect, Release, or file transfer task.
End Time	Date and time when the action executed by the command ends.
# of Files	Number of files included in the action.
Status	Status of the action, such as Success or Failed.

Viewing details of the integration logs

To view details about an integration log, click the instance name on the **File Transfer Logs** page.

Depending on the **Asset Type**, perform one of the following steps to view the details of an integration log:

- To view the file transfer task command logs, click the **Actions** icon in the row that contains the job and select **Download Job Log**. The log is downloaded as a `.txt` file.
  The following image shows the download option for a sample file transfer task command job:

- To view the logs for Collect and Release API commands, click the instance name. To download the log of an instance, click **Download Job Log** in the Results area.
  The following image shows the log of a sample Release command API:

CHAPTER 9

Monitoring elastic clusters

You can monitor the elastic clusters that are running in your organization. An elastic cluster is a set of computing machines that are used to run elastic jobs which are instances of elastic mappings and associated mapping tasks.

When you navigate to the **Elastic Clusters** page in Monitor, you can view a list of all elastic clusters in your organization. You can drill down on each cluster to view details for the cluster on the following tabs:

Activity Log

> A list of cluster events that include the time that a cluster is started, scaled up or down, stopped, and the time that it is modified based on updates to the elastic configuration. It also includes events that make the cluster unusable, as well as cluster recovery events.

Lifecycle Graph

> A visual representation of the number of worker nodes on the cluster over time.

Configuration

> The elastic configuration that a Secure Agent uses to create the elastic cluster. To edit the elastic configuration, you must use Administrator.

Jobs

> A list of all elastic jobs that are running on the cluster. You can stop and restart instances of elastic mappings and associated mapping tasks. You can also download log files.

When you drill down to a cluster instance, the unique cluster ID appears in parentheses in the header of each page. You can use the cluster ID to identify the cluster instance in log files.

Note: The cluster ID in Monitor might not match the cluster ID that appears as the value for the KubernetesCluster tag that is assigned to cloud resources in an AWS environment.

Monitoring all clusters

You can monitor all of the elastic clusters in your organization on the **Elastic Clusters** page.

Use the **Elastic Clusters** page to view all of the elastic clusters that are available to run elastic jobs. For clusters that are currently running, use the page to view a summary of the resource consumption on the cloud platform.

To monitor clusters on the **Elastic Clusters** page, you need the Admin or Designer role. Users with either of these roles can monitor all clusters in an organization. You can also create a custom role and assign the privilege to read elastic configurations to the role.

The following image shows the **Elastic Clusters** page:
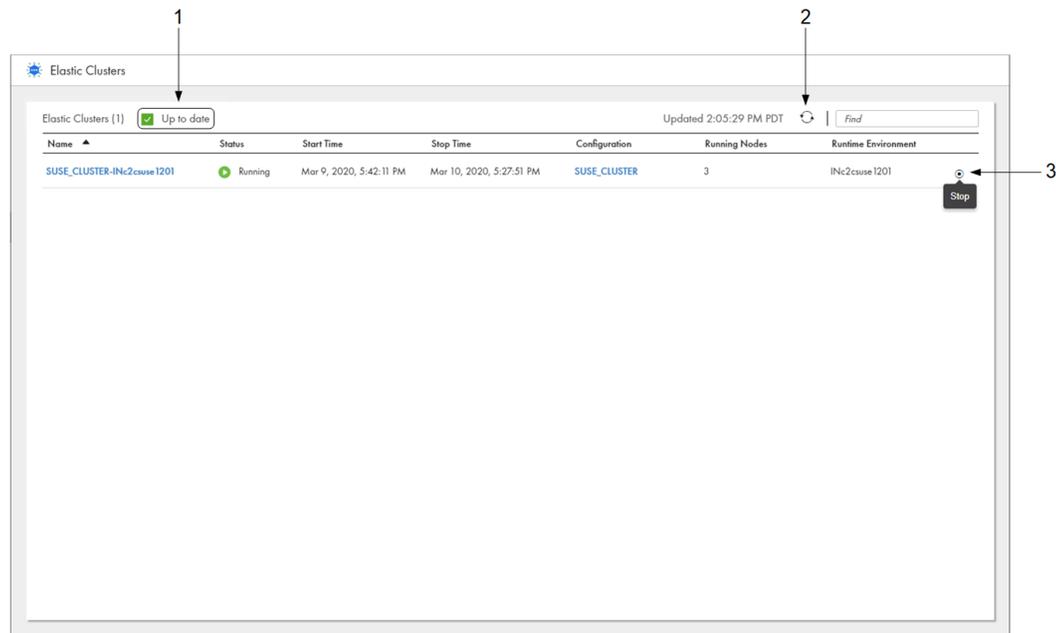
1. Status message that indicates whether information on the page is up-to-date or needs to be refreshed
2. Refresh icon
3. Stop icon

The name of each elastic cluster appears in the following format:

```
<elastic configuration>-<runtime environment>
```

The status message at the top of the page indicates whether the information on the page is up-to-date. If the information is out-of-date, the status message displays "Updates Available." To refresh the page, click the "Updates Available" message or the **Refresh** icon.

Information can become out-of-date when the state of a cluster changes, such as when you run an elastic job to start a cluster or the cluster is scaled up to increase the number of worker nodes.

If a cluster is running, you can stop the cluster by clicking the **Stop** icon. To stop a cluster, you need at least the update privilege for elastic configurations.

You can drill down on a cluster instance to view additional details. To drill down on a cluster instance, click the cluster instance name. You can also drill down directly to the elastic configuration that is used to create the cluster instance. To drill down on the configuration, click the configuration name.

Cluster statuses

For each cluster that you view on the **Elastic Clusters** page, you can view the status of the cluster.

A cluster can be in one of the following statuses:

Status	Description
Starting	The cluster is starting. A cluster starts as soon as you run an elastic job.
Running	The cluster is running. The Serverless Spark engine is processing elastic jobs on the cluster.

Status	Description
Stopping	The cluster is stopping. The elastic jobs that were running on the cluster have completed and the cluster has reached the idle timeout in the elastic configuration, or you recently stopped the cluster in Monitor. The time that it takes to stop a cluster depends on the cloud platform. If you run an elastic job while the cluster is stopping, the cluster does not start and the job fails.
Stopped	The cluster has stopped.
Error	The cluster has an error. During an error, the Secure Agent attempts to recover the cluster. User action might be necessary, such as when you receive a fail-to-start or fail-to-stop exception.
Unknown	The status of the cluster is unknown. If the status is unknown, verify that the Secure Agent is running. If the agent is not running, enable the agent and check whether the cluster starts running. If the cluster does not start running, an administrator can run the command to list clusters. If the command output returns the cluster state as partial or in-use, the administrator can run the command to delete the cluster. For more information about the commands, see *Administrator* in the Administrator help.

Monitor might not reflect the current cluster status if the following conditions are true:

- The Secure Agent machine is shut down.

- You update the elastic configuration, choosing to disable the elastic cluster when you save the configuration.

There is also a delay between the time that the cluster status changes and the time that the agent receives information about the cluster status. So the agent might submit a job to the cluster while the cluster is stopping or stopped. The job fails, and you must run the job again to restart the cluster.

For example, if the agent is notified that the cluster is running and the cluster reaches its idle timeout immediately afterwards, the agent submits the job to the cluster and the job fails.

Monitoring the activity log

You can monitor the activity log for a cluster instance on the **Activity Log** tab after you drill down to a cluster instance from the **Elastic Clusters** page.

Use the activity log to monitor the events on a cluster. The events mark the time that a cluster is started, scaled up or down, stopped, or the time that it is modified based on updates to the elastic configuration. The columns that you can view depend on the cloud platform.

The following image shows the **Activity Log** tab:
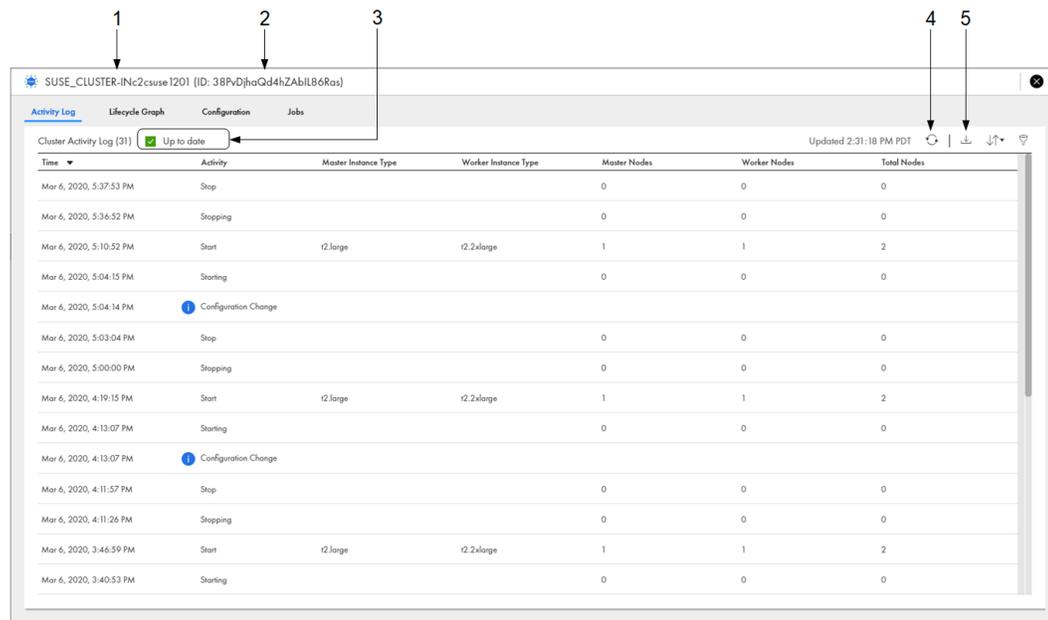
1. Cluster name
2. Cluster ID
3. Status message that indicates whether information on the page is up-to-date or needs to be refreshed
4. Refresh icon
5. Download icon

The status message at the top of the page indicates whether the information on the page is up-to-date. If the information is out-of-date, the status message displays "Updates Available." To refresh the page, click the "Updates Available" message or the **Refresh** icon.

Information can become out-of-date when a new cluster event occurs, such as when you run an elastic job to start the cluster or the cluster is scaled up to increase the number of worker nodes.

To download the activity log, click the **Download** icon.

Cluster events

When you monitor the activity log for a cluster instance, you view a list of cluster events. Events occur on a cluster at a specific point in time.

The following table describes the events that can occur on a cluster:

Cluster event	Description
Starting	The cluster is starting.
Start	The cluster started.
Stopping	The cluster is stopping.
Stop	The cluster stopped.
Scale Up	The number of worker nodes on the cluster increased.

Cluster event	Description
Scale Up Failed	The cluster failed to scale up. The cluster might fail to scale up if an initialization script fails on a worker node that is added to the cluster.
Scale Down	The number of worker nodes on the cluster decreased.
Configuration Change	The elastic configuration was changed. The cluster is stopped at the time of a configuration change. The changes in the configuration take effect the next time that the cluster starts.
Unusable	The cluster entered an error status.
Recovery	The cluster was recovered after encountering an error.

Viewing the lifecycle graph

You can view the lifecycle graph for a cluster instance on the **Lifecycle Graph** tab after you drill down to a cluster instance from the **Elastic Clusters** page.

The lifecycle graph is a visual representation of the number of worker nodes on the cluster over time. You can change the time range to view more or less granulated details about the changes to the number of worker nodes.

The following image shows the **Lifecycle Graph** tab:
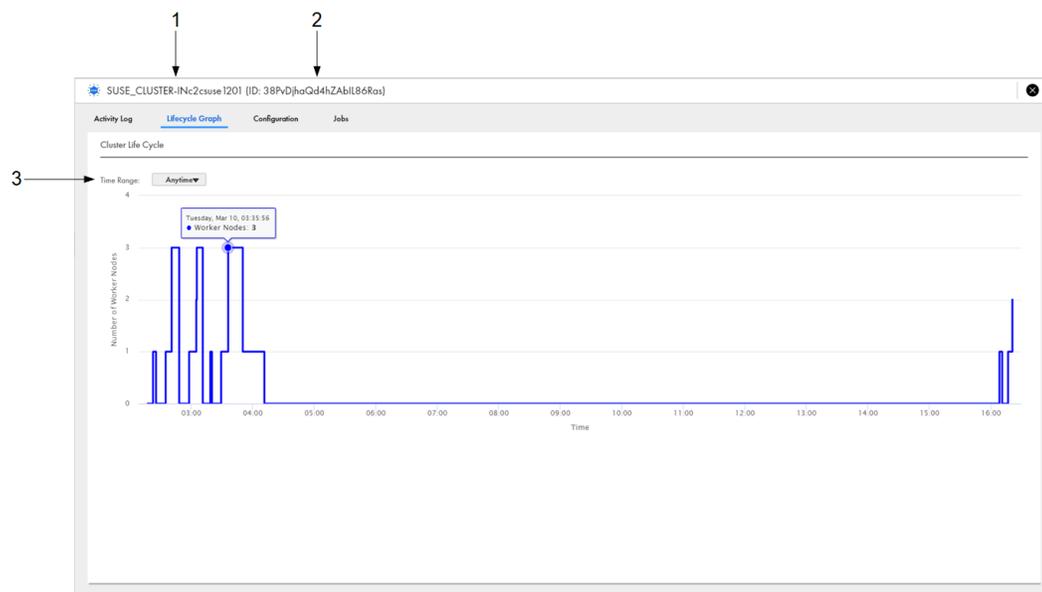
1. Cluster name
2. Cluster ID
3. Time range

Viewing the configuration

You can view the configuration for a cluster instance on the **Configuration** tab after you drill down to a cluster instance from the **Elastic Clusters** page. The configuration that you view is the elastic configuration that you use to provision resources for an elastic cluster.

Use the **Configuration** tab to reference the configuration. The properties that you can view depend on the cloud platform. To edit the configuration, use Administrator.

The following image shows the **Configuration** tab:
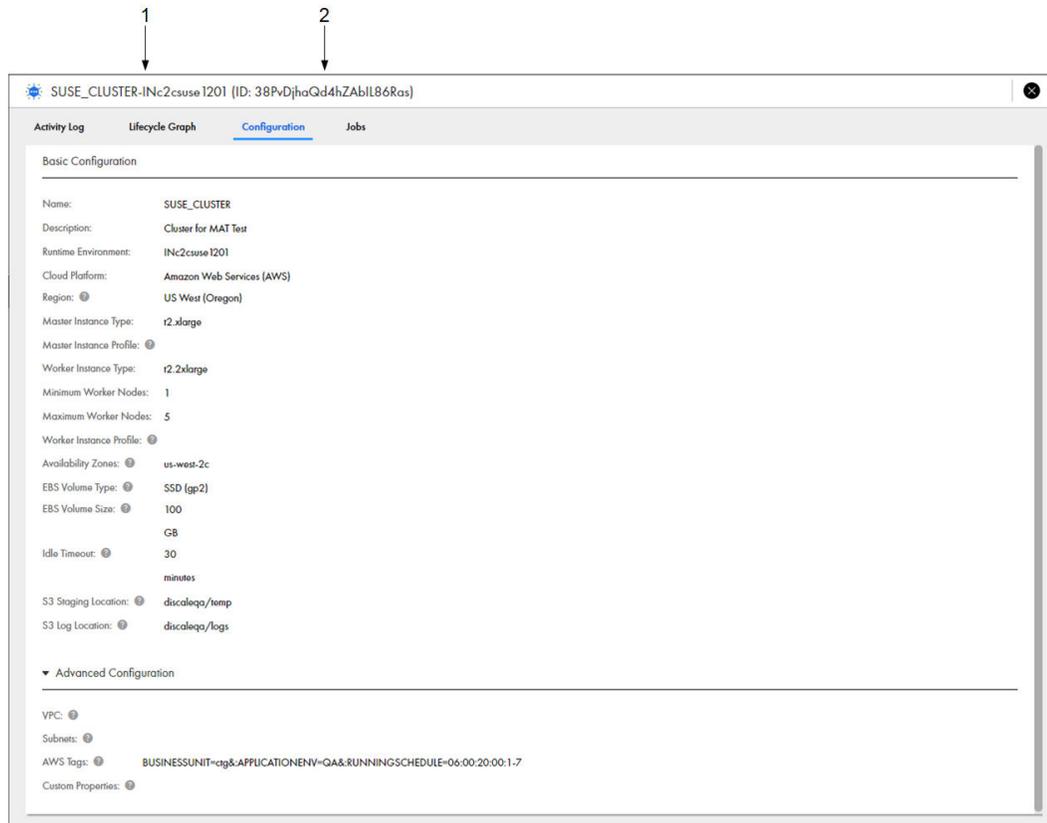
1. Cluster name
2. Cluster ID

Monitoring jobs on a cluster

You can monitor all elastic jobs that are running on a cluster on the **Jobs** tab after you drill down to the cluster instance from the **Elastic Clusters** page. The **Jobs** tab lists the jobs that are currently running and the jobs that have completed.

Use the **Jobs** tab for failure analysis and debugging of the elastic jobs or the elastic cluster.

To avoid unnecessary failures, check the status of the elastic cluster before you run an elastic job on the cluster. The cluster should either not exist, be running, or be stopped.

The following image shows the **Jobs** tab:
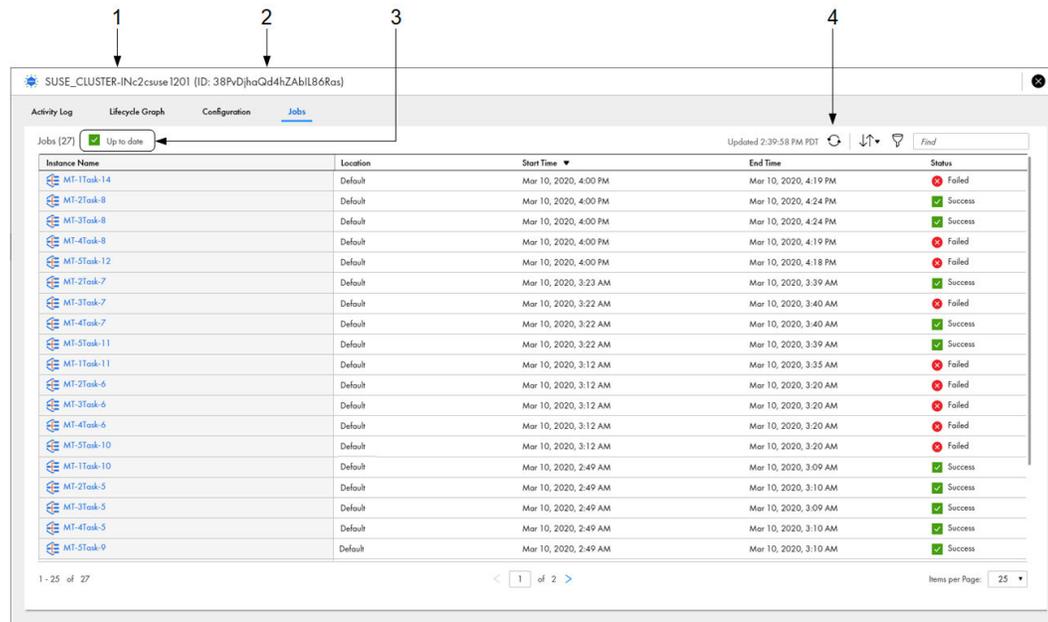
1. Cluster name
2. Cluster ID
3. Status message that indicates whether information on the page is up-to-date or needs to be refreshed
4. Refresh icon

The **Jobs** tab lists the jobs that were run within the last three days, plus the 1000 most recent jobs that are more than three days old.

The status message at the top of the page indicates whether the information on the page is up-to-date. If the information is out-of-date, the status message displays "Updates Available." To refresh the page, click the "Updates Available" message or the **Refresh** icon.

Information can become out-of-date when a job status changes or when a user starts a job.

When a job completes, you can drill down on the job to view the job details. To drill down on a job, click the instance name.

For information about the job details, see "Monitoring elastic mappings and tasks" on page 21.

CHAPTER 10

Monitoring source control logs

If your organization is enabled for source control, you can monitor the actions users in your organization perform on source controlled objects. Each time a user performs an action on an object that is source controlled, Informatica Intelligent Cloud Services logs the action on the **Source Control Logs** page. Logs are retained for seven days.

To monitor source control actions you need the **Asset - Source Control Logs** privilege. Users with the following roles have the **Asset - Source Control Logs** privilege by default:

- Administrator
- Designer
- Monitor

The following image shows the **Source Control Logs** page:
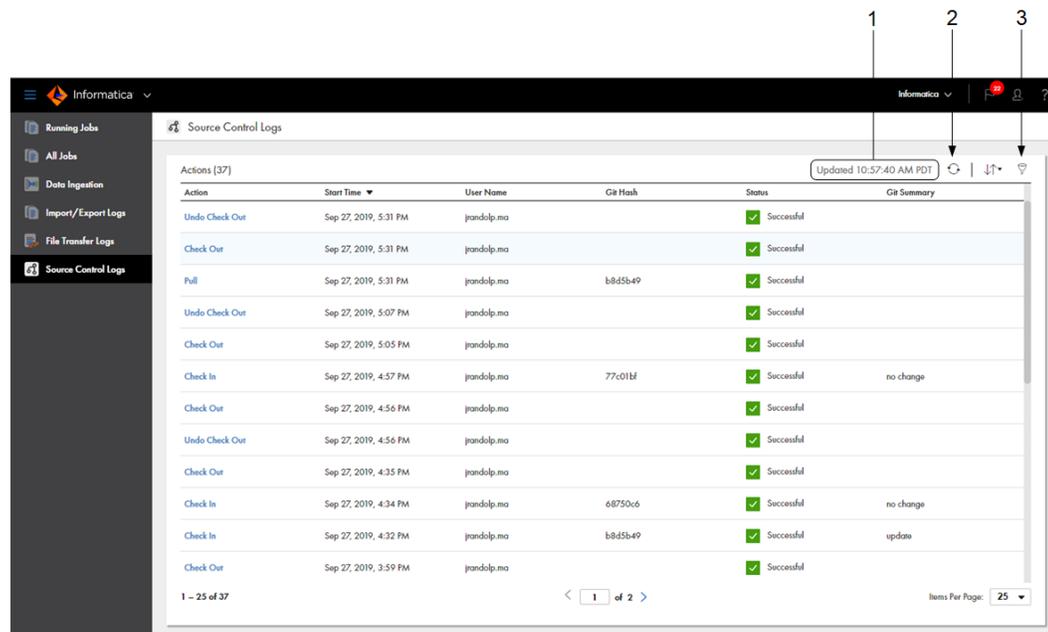
1. Time the page was last updated.
2. Refresh icon.
3. Filter icon.

By default, the **Source Control Logs** page lists actions beginning with the most recent. To reorder the page, click the column name of the property you want to sort by. You can sort actions by the following properties:

- Action

- Start Time
- User Name
- Status

To refresh the page, click **Refresh**.

To filter the actions that appear on the **Source Control Logs page**, click the **Filter** icon. You can specify keywords and partial strings in your filters.

To apply a filter, click **Add Filter**, select the property to filter by, and then enter the property value. For example, to find all the check ins with a Git hash that contain the string "b8d", add the **Action** filter and select **Check In**. Then add the **Git Hash** filter and enter "b8d" as the value.

To view detailed information about an action, click the action name.

Action properties

The **Source Control Logs** page displays action properties such as the name, start time, and git hash.

The action properties include the following properties:

Property	Description
Action	Git action that was performed.
Start Time	Date and time the action was started.
User Name	User who performed the action.
Git Hash	Git identifier for the revision. Informatica Intelligent Cloud Services displays the first seven characters of the git hash. Applicable for the following actions: - Check in - Pull
Status	Action status. An action can have one of the following statuses: - Success. The action completed successfully. - In progress. The action is currently in progress. - Warning. The action completed with warnings. See the asset summary for more information. - Error. The action did not complete because of an error.
Git Summary	Description of the revision. Applicable for the following actions: - Check in - Pull

Viewing details for a source control action

To view detailed information about a source control action, click the action name in the Action column.

The following image shows the details for a check in:
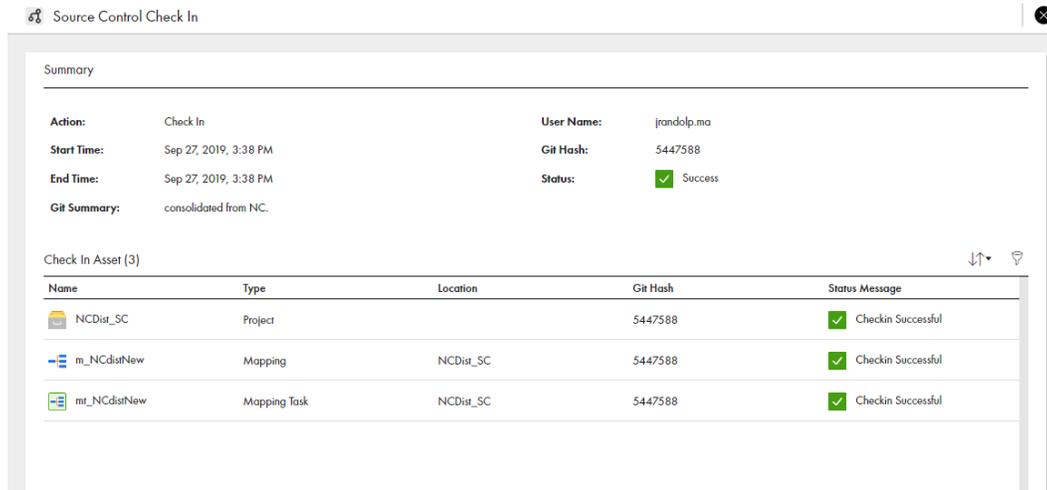
Action summary

The **Summary** area displays general details about the action.

The action summary includes the following properties:

Property	Description
Action	Action that was performed in the source control repository.
Start Time	Date and time the action was started.
End Time	Date and time the action ended.
Git Summary	Description of the revision. Applicable for the following actions: - Check in - Pull
User Name	User who performed the action.
Git Hash	Git identifier for the revision. Informatica Intelligent Cloud Services displays the first seven characters of the git hash. Applicable for the following actions: - Check in - Pull
Status	Status of the action. An action can have one of the following statuses: - Success. Action was completed successfully. - In Progress. The action is in progress. - Warning. The action completed with warnings. See the asset summary for more information. - Error. Action was not perform because of an error.

Asset summary

The **Assets** area displays the project, folder, and assets the action was performed on and details about each object.

The asset summary includes the following properties for each object:

Property	Description
Name	Name of the project, folder, or asset.
Type	Type of object.
Location	Project or folder in which the asset resides.
Git Hash	Git identifier for the action. Applicable for the following actions: - Check in - Pull
Status Message	Status of the individual object and details about the status.

INDEX