



Informatica® Intelligent Cloud Services
August 2024

Monitor

Informatica Intelligent Cloud Services Monitor
August 2024

© Copyright Informatica LLC 2006, 2024

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

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

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

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

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

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

Publication Date: 2024-08-21

Table of Contents

Preface	7
Informatica Resources.	7
Informatica Documentation.	7
Informatica Intelligent Cloud Services web site.	7
Informatica Intelligent Cloud Services Communities.	7
Informatica Intelligent Cloud Services Marketplace.	7
Data Integration connector documentation.	8
Informatica Knowledge Base.	8
Informatica Intelligent Cloud Services Trust Center.	8
Informatica Global Customer Support.	8
Chapter 1: Monitoring jobs.....	9
Monitoring all jobs.	9
Monitoring running jobs.	10
Monitoring your jobs.	11
Monitoring subtasks.	12
Job properties.	14
Customizing the jobs pages.	16
Viewing details for a specific job.	17
Chapter 2: Monitoring Data Integration jobs.....	19
Monitoring mappings and mapping tasks.	20
Job properties.	20
Job results.	21
Individual source and target results.	22
Parameter details.	22
Sequence generator details.	22
SQL ELT optimization details.	23
Monitoring mappings and tasks in advanced mode.	23
Job properties.	23
Job results.	24
Individual task results.	25
Runtime plan.	26
Reprocessing job details.	26
Incremental file load details.	26
CLAIRE recommendations.	27
Monitoring advanced cluster subtasks.	28
Job properties.	28
Job results.	29
Job results for tuning.	29

Spark task details.	30
Monitoring code tasks.	31
Job properties.	31
Job results.	32
Code task API execution parameters.	33
Spark application task details.	34
Monitoring dynamic mapping tasks.	35
Job properties.	35
Job results.	36
Individual job results.	36
Monitoring synchronization tasks.	37
Job properties.	37
Job results.	37
Monitoring data transfer tasks.	38
Job properties.	38
Job results.	39
Individual source and target results.	39
Monitoring masking tasks.	40
Job properties.	40
Job results.	40
Activity log details.	41
Monitoring replication tasks.	41
Job properties.	42
Job results.	42
Individual object results.	43
Monitoring PowerCenter tasks.	43
Job properties.	43
Job results.	44
Monitoring taskflows.	45
Start step properties.	47
Data Task properties.	47
IntegrationOps Task properties.	48
Notification Task properties.	49
Command Task step properties.	49
File Watch Task step properties.	50
Ingestion Task step properties.	52
Subtaskflow step properties.	53
Throw step properties.	53
End step properties.	53
Monitoring linear taskflows.	54
Job properties.	54
Job results.	55

Individual task results.	55
Running only a particular step in a taskflow instance.	56
Running a taskflow instance from a step.	57
Updating variables to run a taskflow instance from a step.	58
Stopping and restarting mapping and task instances.	59
Stopping, suspending, resuming, and restarting taskflow instances.	60
Taskflow log files.	61
Downloading a taskflow log file from Monitor.	61
Downloading a taskflow log file using the log resource.	61
Taskflow log file contents	62
Sample taskflow log file	63
Chapter 3: Data Integration job log files.	65
Viewing an error rows file.	67
Downloading a session log file.	67
Viewing log files for advanced cluster subtasks.	68
Chapter 4: Monitoring Data Profiling jobs.	70
Job properties.	70
Subtask types.	71
Stopping and restarting data profiling jobs.	72
Downloading a log file.	73
Chapter 5: Monitoring imports and exports.	74
Import and export log properties.	74
Customizing import and export logs.	75
Viewing details for an import or export.	76
Downloading an export file.	77
Downloading an import or export log.	78
Chapter 6: Monitoring file transfer jobs.	79
Monitoring AS2 file transfers.	79
Viewing AS2 file transfer logs.	79
Viewing details for an AS2 file transfer.	80
Monitoring SFTP file transfers.	83
Viewing SFTP file transfer logs.	84
Viewing details of an SFTP file transfer log.	85
Monitoring HTTPS file transfers.	86
Viewing HTTPS file transfer logs.	86
Viewing details of HTTPS file transfer log.	87
Monitoring MLLP file transfers.	87
Viewing MLLP file transfer logs.	88
Viewing details of an MLLP file transfer log.	89

Monitoring file listeners.	89
File listener job details.	90
Monitoring integration APIs	91
Viewing integration logs.	92
Viewing details of the integration logs.	93
Chapter 7: Monitoring Data Ingestion and Replication jobs.....	96
Monitoring your ingestion and replication jobs.	97
Monitoring all ingestion and replication jobs.	98
Job properties.	101
Viewing details for an ingestion and replication job.	103
Application ingestion and replication job details.	103
Database ingestion and replication job details.	109
File ingestion and replication job details.	116
Streaming ingestion and replication job details.	119
Data Ingestion and Replication alerts.	121
Configuring alerts for Data Ingestion and Replication jobs.	122
Chapter 8: Monitoring advanced clusters.....	123
Monitoring an individual advanced cluster.	125
Cluster statuses.	126
Monitoring an advanced cluster instance.	126
Monitoring the activity log.	127
Viewing the lifecycle graph.	129
Viewing the configuration.	130
Monitoring jobs on a cluster.	130
Chapter 9: Monitoring source control logs.....	132
Action properties.	133
Viewing details for a source control action.	134
Action summary.	134
Asset summary.	135
Index.....	136

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 ServicesSM 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, on the [Informatica Intelligent Cloud Services Status](#) page, click **SUBSCRIBE TO UPDATES**. You can choose to receive notifications sent as emails, SMS text messages, webhooks, RSS feeds, or any combination of the four.

Informatica Global Customer Support

You can contact a Global Support Center through the Informatica Network or by telephone.

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

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

CHAPTER 1

Monitoring 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.

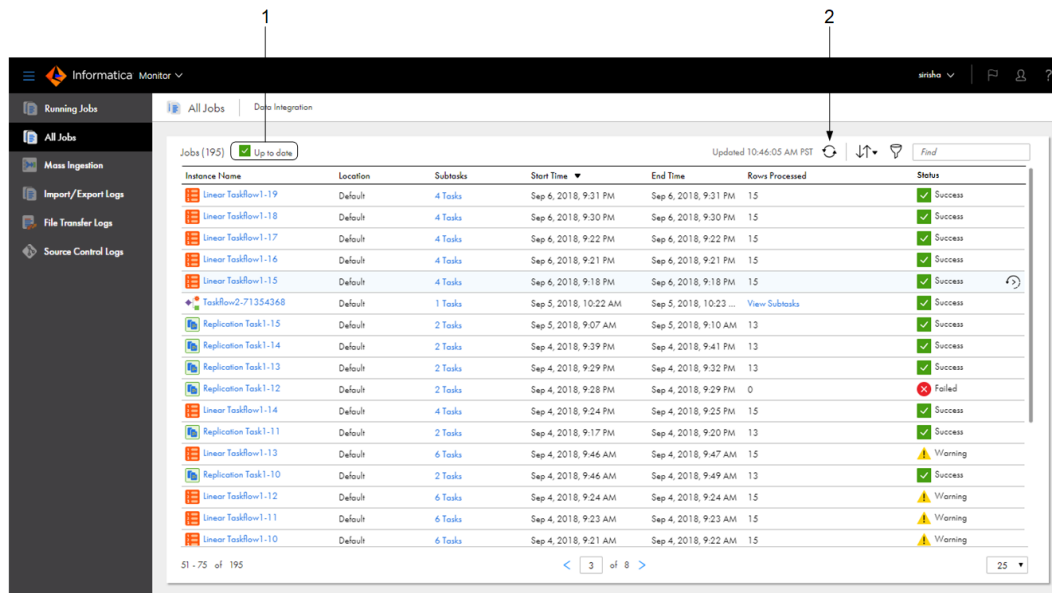
Use the **All Jobs** page for failure analysis and debugging of the jobs in your organization.

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 all the jobs that were run within the last three days. If your organization has run less than 1000 jobs, the page also includes older jobs, up to 1000 total jobs.

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:

Instance Name	Location	Subtasks	Start Time	End Time	Rows Processed	Status
MappingTask1-2	Default		Feb 7, 2019, 2:18 PM		0	Running
MappingTask4-1	Default		Feb 7, 2019, 2:17 PM	Feb 7, 2019, 2:17 PM	0	Success
Mapping6-2	Default		Feb 7, 2019, 2:16 PM	Feb 7, 2019, 2:16 PM	1999	Success
Mapping5-2	Default		Feb 7, 2019, 2:15 PM	Feb 7, 2019, 2:15 PM	0	Success
Synchronization Task4-relatedObj-1	Accounts\February2018		Feb 7, 2019, 2:15 PM	Feb 7, 2019, 2:15 PM	0	Failed
MappingTask3-1	Default		Feb 7, 2019, 2:15 PM	Feb 7, 2019, 2:15 PM	0	Success
Mapping2-2	Default		Feb 7, 2019, 2:14 PM	Feb 7, 2019, 2:15 PM	0	Success

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.

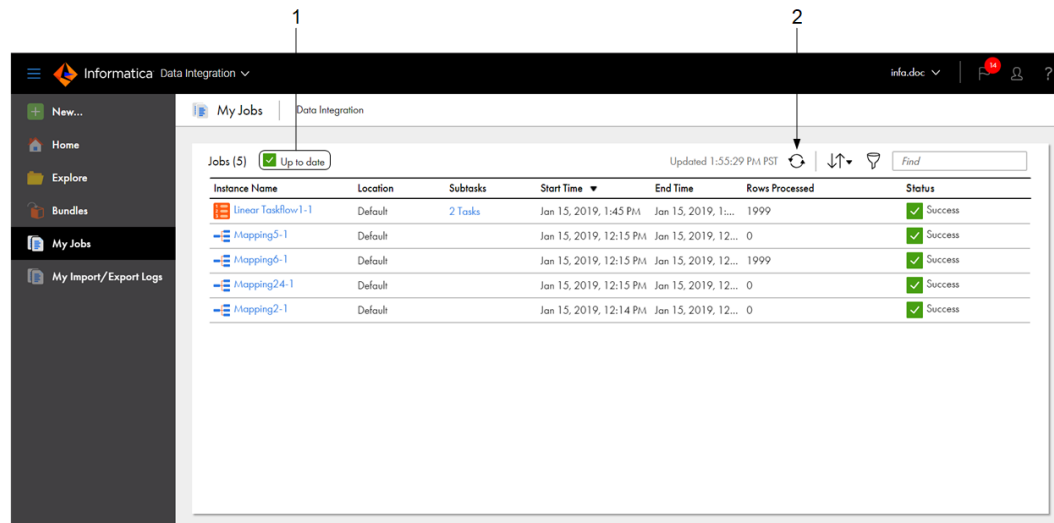
Monitoring your jobs

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
- Data Ingestion and Replication

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 all the jobs that you ran within the last three days. If you have run less than 1000 jobs in the last three days, the log also includes older jobs, up to 1000 total jobs.

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 task instance that contains subtasks, you can also monitor the subtasks. When you monitor a taskflow instance that contains subtasks or subtaskflows, you can also monitor the subtasks and subtaskflows. You can monitor a subtask or subtaskflow 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 or subtaskflows, 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:

- Mappings in advanced mode
- Replication tasks
- Advanced taskflows
- Linear taskflows

- Azure data sync tasks
- Data Profiling tasks
- Dynamic mapping tasks
- File listener
- File ingestion task

The following image shows the **My Jobs** page in Data Integration when you view subtasks for a linear taskflow:



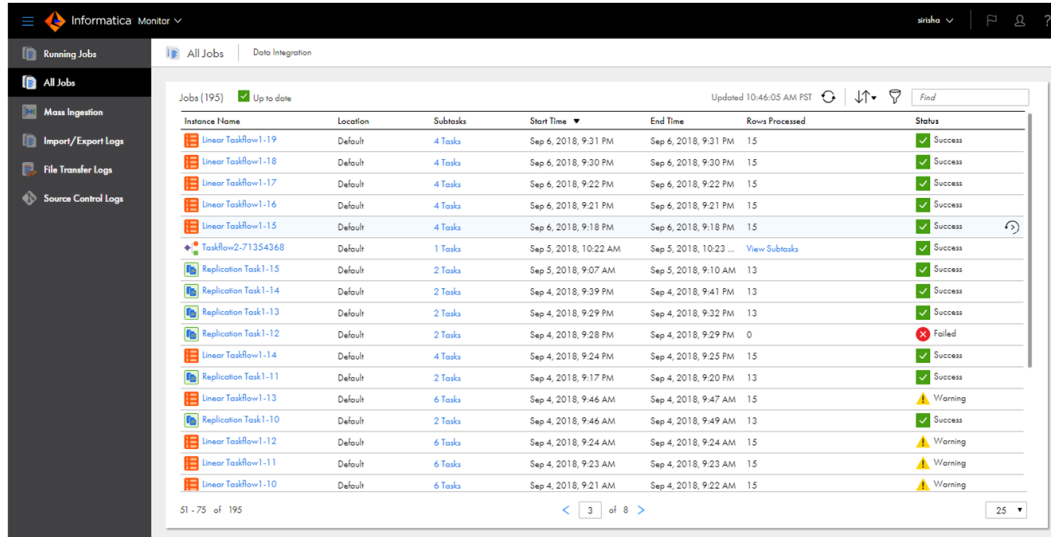
1. Click the service name to close the subtasks view.

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	<p>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.</p> <p>If the job you are viewing is a non-linear taskflow instance, this field displays View Subtasks. Click View Subtasks to monitor the subtasks and subtaskflows for the taskflow. If the taskflow contains nested subtasks and subtaskflows, this field displays View Subtasks to view the nested subtasks and subtaskflows.</p> <p>Does not apply to advanced cluster subtasks.</p>
Status	<p>Job status. A job can have one of the following statuses:</p> <ul style="list-style-type: none"> - 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 job was stopped by the user. Does not apply to data transfer tasks or dynamic mapping tasks. <p>For subtasks of replication task instances, the parent job has stopped running, so the subtask cannot start.</p> <ul style="list-style-type: none"> - Stopping. The job is stopping. Applies to tasks that are stopped cleanly. - 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 file ingestion and replication task instances. - Failed. The job did not complete because it encountered errors. <p>Tip: If the job status is Failed, you can hover the cursor over the job status to view and copy the error message.</p>


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 file ingestion and replication tasks or advanced cluster subtasks.
Failure Rows	Number of rows that were not written to the target. Does not apply to file ingestion and replication tasks or advanced cluster subtasks.
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 screenshot displays the Informatica Data Integration interface. The top navigation bar includes the Informatica logo, 'Data Integration', and a dropdown menu for 'info-stage1'. A sidebar on the left contains navigation options: 'New...', 'Home', 'Explore', 'Bundles', 'My Jobs', 'My Import/Export Logs', and the selected job 'lt_Cust-2'. The main content area shows the job details for 'lt_Cust-2'.

Job Properties

- Task Name: [lt_Cust](#)
- Instance ID: 2
- Task Type: Linear Taskflow
- Started By: Itroy through UI
- Start Time: Aug 14, 2018 3:47:33 PM
- End Time: Aug 14, 2018 3:47:42 PM
- Duration: 9 seconds

Results

- State: ✔ Success
- Success Rows: 2761
- Errors: 0

Individual Task Results

#	Task Name	End Time	Status	Success Rows	Errors	Error Message	Recommendation
1	mt_Cust	Aug 14, 2018 3:47:37 PM	✔ Success	2499	0		
	src_AllCust			2499	0		
	AllCustomers_OUT_csv			2499	0		
2	mt_CustOrders	Aug 14, 2018 3:47:42 PM	✔ Success	262	0		
	src_CustOrders			262	0		
	CustOrders_OUT_csv			262	0		

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.

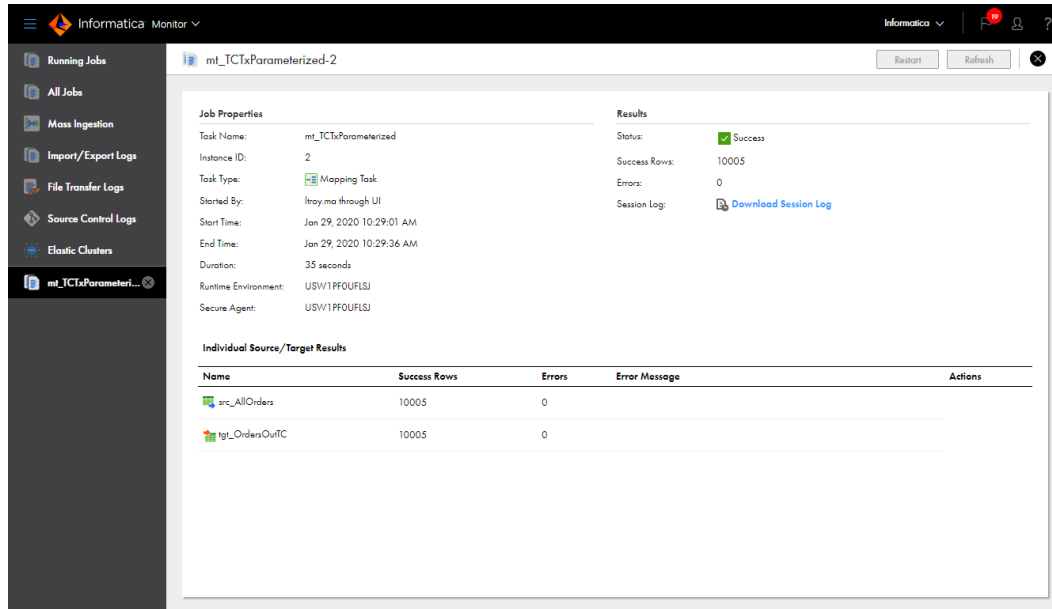
You can monitor the following types of Data Integration jobs:

- Mappings
- Mapping data preview jobs
- Mapping tasks
- Mappings in advanced mode
- Subtasks that run on the Data Integration Server
- Subtasks that run on an advanced cluster
- Dynamic mapping tasks
- Synchronization tasks
- Data transfer tasks
- Masking tasks
- Replication tasks
- File ingestion and replication tasks
- PowerCenter 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	<p>Job status. A job can have one of the following statuses:</p> <ul style="list-style-type: none"> - Starting. The job is starting. - Running. The job is still running. - Success. The job completed successfully. - Stopped. The job was stopped by a user or the parent job has stopped running, so the subtask can't 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.
Session Log	<p>Allows you to download the session log file. By default, Informatica Intelligent Cloud Services stores session logs for 10 runs before it overwrites the logs with the latest runs. If you want to retain session logs for more than 10 runs, you can configure the Maximum Number of Log Files property in the task wizard.</p> <p>Session log files are written to the following directory:</p> <pre><Secure Agent installation directory>/apps/Data_Integration_Server/logs</pre>
Requested Compute Units Per Hour	<p>Number of serverless compute units per hour that the task requested.</p> <p>You can view the number of requested compute units if the task runs in a serverless runtime environment.</p>
Total Consumed Compute Units	<p>Total number of serverless compute units that the task consumed.</p> <p>You can view the number of consumed compute units if the task runs in a serverless runtime environment.</p>
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 mappings in advanced mode. 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 transformation or name of the target transformation and the target object.
Success rows	Number of rows successfully written to the target.
Affected rows	Number of rows affected by the defined task operation. Doesn't apply to jobs based on mappings configured for full SQL ELT optimization or mappings in SQL ELT mode.
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.
Actions	Actions that you can take regarding the job.

Parameter details

Parameter details are displayed for job instances that contain in-out parameters or connection or object parameters that were overridden at runtime.

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

For mapping task instances that include connection or object parameters that were overridden at runtime, the job details also display the location of the parameter file and the value of each parameter after the task has completed.

Sequence generator details

Sequence generator details are shown for mapping task and mapping data preview job instances that contain a Sequence Generator transformation that does not use a shared sequence.

The following table describes the properties shown for each sequence generator:

Property	Description
Transformation Name	Name of the Sequence Generator transformation.
Next Value	Value of the NEXTVAL output field in the Sequence Generator transformation.

SQL ELT optimization details

SQL ELT optimization details are shown for mapping task instances where the task is configured for SQL ELT optimization.

Expand the **SQL ELT Optimization** area to view SQL ELT details.

The **SQL ELT Optimization** area displays the SQL ELT type configured in the task and the type of SQL ELT the task executed. The area also displays the SQL queries pushed to the source or target, along with any errors.

Monitoring mappings and tasks in advanced mode

To view detailed information about a mapping in advanced mode, such as a mapping instance, task instance, or 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 in advanced mode:

The screenshot displays the job details for 'MCT_UnifiedEngineFunctions-1'. The 'Job Properties' section includes:

- Task Name: MCT_UnifiedEngineFunctions
- Instance ID: 1
- Task Type: Mapping Task
- Started By: unified through API
- Start Time: Sep 12, 2023, 6:30:35 PM
- End Time: Sep 12, 2023, 6:36:24 PM
- Duration (HH:MM:SS): 00:05:49
- Runtime Environment: informatica.com
- Job Priority: Medium
- Target Duration Status: Met

The 'Results' section shows a 'Success' status with a 'Download' link for the session log. Below this is a table of 'Individual Task Results':

Name	Runtime	Start time	End time	Success Rows	Error Rows	Error Message	Status
InfoSpark0	Advanced Cluster	Sep 12, 2023, 6:32:50 PM	Sep 12, 2023, 6:34:02 PM				Success
CDITask	Data Integration Server	Sep 12, 2023, 6:34:55 PM	Sep 12, 2023, 6:36:09 PM				Success

Job properties

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

The following table describes the job 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, such as Mapping (Advanced Mode).

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.
Job Priority	Priority used to schedule the job.
Target Duration Status	Indicates whether the job met or exceeded its target duration.

Job results

The job results display the status of the job, a download link to the session log, and an error message for each mapping instance, task instance, or data preview job in advanced mode.

The following table describes the job results:

Property	Description
Status	<p>Job status. A job can have one of the following statuses:</p> <ul style="list-style-type: none"> - Starting. The job is starting. - Running. The job is still running. - Stopped. The job was stopped by a user or the parent job has stopped running, so the subtask can't start. - Success. The job completed successfully. - Warning. The job completed with errors. - Failed. The job did not complete because it encountered errors.
Session Log	<p>Allows you to download the session log file. By default, Informatica Intelligent Cloud Services stores session logs for 10 runs before it overwrites the logs with the latest runs. If you want to retain session logs for more than 10 runs, you can configure the Maximum Number of Log Files property in the task wizard.</p> <p>Session log files are written to the following directory:</p> <pre><Secure Agent installation directory>/apps/Data_Integration_Server/logs</pre>
Requested Compute Units Per Hour	<p>Number of serverless compute units per hour that the task requested.</p> <p>You can view the number of requested compute units if the task runs in a serverless runtime environment.</p>
Total Consumed Compute Units	<p>Total number of serverless compute units that the task consumed.</p> <p>You can view the number of consumed compute units if the task runs in a serverless runtime environment.</p>
Error Message	Error message, if any, that is associated with the job.

Individual task results

The individual task results display results for subtasks that process the mapping in advanced mode.

Each subtask runs either on the Data Integration Server or on an advanced cluster:

- **Data Integration Server.** You can expand the Data Integration Server subtask to view the source and target results.
- **Advanced cluster.** You can drill down on the advanced cluster subtask. For more information, see [“Monitoring advanced cluster subtasks” on page 28](#).

You can hover over each row in the individual task results to download a detailed log for Data Integration Server subtasks or to download the agent and driver logs for advanced cluster subtasks.

The following table describes the individual task results:

Property	Description
Name	Name of the subtask.
Runtime	Indicates whether the subtask ran on the Data Integration Server or on an advanced cluster.
Start time	Date and time that the subtask started.
End time	Date and time that the subtask completed or stopped.
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. For advanced cluster subtasks, this field displays a value only if the task runs using SQL ELT optimization. For Spark subtasks, this field displays a value only for supported connectors.
Error Rows	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. For advanced cluster subtasks, this field displays a value only if the task runs using SQL ELT optimization. For Spark subtasks, this field displays a value only for supported connectors.
Error Message	Error message, if any, that is associated with the subtask.
Status	Status of the subtask. A subtask can have one of the following statuses: <ul style="list-style-type: none">- Queued. The subtask is queued on a Secure Agent, but it has not started yet.- Starting. The subtask is starting.- Running. The subtask is still running.- Success. The subtask completed successfully.- Stopped. The parent task stopped running, so the subtask cannot start.- Warning. The subtask completed with errors.- Failed. The subtask did not complete because it encountered errors.

Runtime plan

The runtime plan is a visualization that closely approximates how the data actually flows through the mapping at runtime. You can view which mapping logic runs on the Data Integration Server or on an advanced cluster.

Note: If Monitor fails to retrieve the runtime plan, download the session log and review the mapping compilation log in the archive file. You can search for the following line: `[LDTM_7343] Generating runtime plan.`


For more information about runtime plans, see *Mappings* in the Data Integration help.

Reprocessing job details

If you run a job to reprocess incrementally-loaded source files, you can view details about the reprocessing job.

You can view the start and end times of the time range that the job uses to search for new and modified files in the configured source directories. The start time is the time that you configure in the advanced options. If you configure reprocessing in a time interval, the end time is time that you configure in the advanced options. Otherwise, the time at which the job runs determines the end time.

The following image shows an example of the monitoring details for a reprocessing job:



Note: If you restart a job that reprocessed incrementally-loaded source files, the new job runs as a regular job without the advanced options.

Incremental file load details

If you configure sources to incrementally load files, you can view details about the file load.

You can view the start and end times of the time range that the job uses to search for new and modified files in a directory.

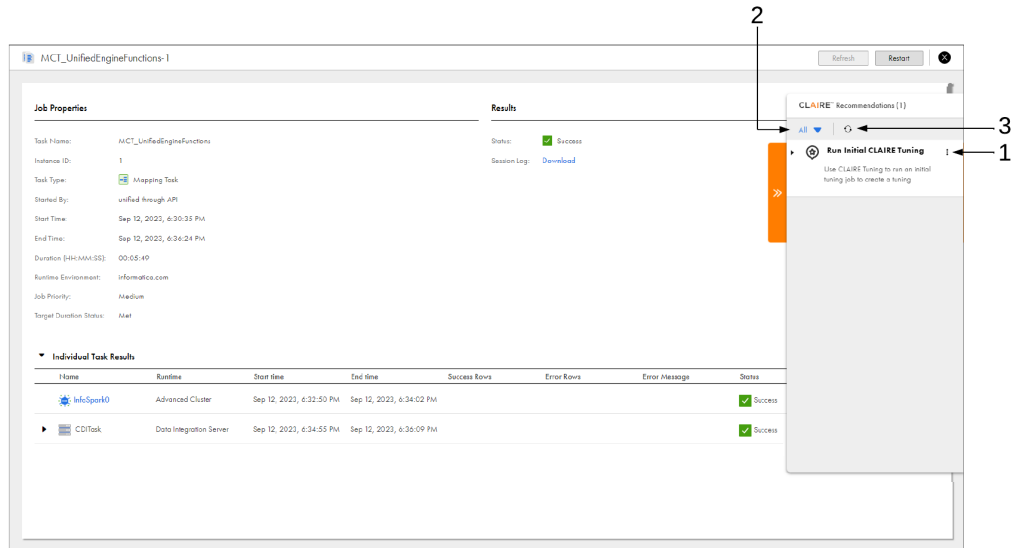
The start load time is derived from the end load time of the previous job. The start load time does not apply when the job runs for the first time or you reset the last load time in the mapping task.

Jobs process all files modified up to the end load time, including up to 999 milliseconds after the second. For example, if the end load time of a job appears as Aug 28, 2019 2:15:22 p.m., the job processes all files modified up to Aug 28, 2019 2:15:22.999 p.m.

CLAIRE recommendations

View CLAIRE recommendations to improve job performance and to reduce job costs in the **Recommendations** panel. CLAIRE recommendations are available if CLAIRE is enabled in your organization.

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



1. Actions menu

Use the **Actions** menu to mark recommendations as complete or incomplete, or to opt in and opt out of recommendations.

CLAIRE automatically applies some recommendations. You can use the **Actions** menu to opt out of the recommendation or opt back in.

Other recommendations require manual action. These recommendations appear as to-do items. You can use the **Actions** menu to mark to-do items as complete, or mark them as incomplete. You can also opt out of the recommendation or opt back in.

2. Filter recommendations

Use the **Filter** menu to filter recommendations. You can use the following filters:

- **All** shows all recommendations.
- **To Do** shows recommendations that require manual action.
- **Applied** shows recommendations that are applied automatically and recommendations that are marked complete.
- **Opted-Out** shows recommendations that are opted out of.

3. Refresh recommendations

Refresh the recommendations to update the recommendations in the **Recommendations** panel.

Monitoring advanced cluster subtasks

To view detailed information about a subtask that runs on an advanced cluster, navigate to the subtask from the individual task results for a mapping in advanced mode.

Note: When you monitor a subtask that runs on an advanced cluster 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 subtask that runs on an advanced cluster display general properties about the instance.

The following table describes the job properties for subtasks that run on an advanced cluster:

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.
Advanced configuration	Advanced configuration that was used to create the advanced cluster.
Cluster	Advanced cluster where the 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 subtask that runs on an advanced cluster display the status of the job, a download link to the Spark execution plan, and an error message, if any.

The job results include the following properties:

Property	Description
Status	<p>Job status. A job can have one of the following statuses:</p> <ul style="list-style-type: none">- Starting. The job is starting.- Running. The job is either queued or running.- Success. The job completed successfully.- Failed. The job did not complete because it encountered errors. <p>If the advanced cluster is not running when you run a job, the job waits for the cluster to start. During this time the job status is Starting.</p> <p>If the Secure Agent fails while the job is running, the status of the job continues to display Running. You must cancel the job and run the job again.</p> <p>The status for a queued job displays Running. To find out if a job is queued or running, check the session log.</p>
Execution Plan	<p>Allows you to download the Spark execution plan which shows the runtime Scala code that the advanced cluster uses to run the data logic in the mapping. You can use the Scala code to debug issues in the mapping.</p>
Error Message	<p>Error message, if any, that is associated with the job.</p>

Job results for tuning

If you tune a mapping task that runs on an advanced cluster, the job results show the tuning job status and a link to the subtasks.

The job results include the following properties:

Property	Description
Status	<p>Job status. A job can have one of the following statuses:</p> <ul style="list-style-type: none">- Running. The job is either queued or running.- Success. The job completed successfully.- Stopped. The job was stopped.- Failed. The job did not complete because it encountered errors. <p>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.</p> <p>The status for a queued job displays Running. To find out if a job is queued or running, check the session log.</p>
Subtasks	<p>Number of subtasks that are part of the tuning job. Each subtask represents a run of the mapping task.</p> <p>If a link is available, click the link to monitor each mapping task.</p>
Requested Compute Units Per Hour	<p>Number of serverless compute units per hour that the task requested.</p> <p>You can view the number of requested compute units if the task runs in a serverless runtime environment.</p>

Property	Description
Total Consumed Compute Units	Total number of serverless compute units that the task consumed. You can view the number of consumed compute units if the task runs in a serverless runtime environment.
Error Message	Error message, if any, that is associated with the job.

Spark task details

In subtasks that run on an advanced cluster, the mapping is translated into Spark tasks that process the data logic simultaneously. You can view details for each Spark task listed under **<Spark task name> Task Results**.

The following table describes the details for each Spark task:

Property	Description
Status	Status of the Spark task. The Spark task can have one of the following statuses: <ul style="list-style-type: none"> - 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 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.
Duration	Amount of time that the Spark task ran.
Memory Per Executor	Amount of memory that each Spark executor uses.
Cores Per Executor	Number of cores that each Spark executor uses.
Driver and Agent Job Logs	Select Download to download the Spark driver and agent job logs.
Advanced Log Location	The log location that is configured in the advanced configuration for the advanced cluster. You can navigate to the advanced log location to view and download the agent job log, Spark driver log, and Spark executor logs.

Each Spark task is translated into Spark jobs, which are further broken down into stages. You can view the following details for each Spark job and stage:

Property	Description
Job Name	Name of the Spark job or stage.
Start time	Date and time when the Spark job or stage started. Start time might be "NA" for aborted tasks.

Property	Description
End time	Date and time when the Spark job or stage ended. End time might be "NA" for aborted tasks.
Duration	Amount of time that the Spark job or stage ran.
Total Tasks	Number of tasks the Spark job or stage attempted.
Successful Tasks	Number of tasks the Spark job or stage successfully completed.
Failed Tasks	Number of tasks that the Spark job or stage failed to complete.
Running Tasks	Number of tasks that the Spark job or stage is currently running.
Input Size / Records	Size of the file and number of records input by the Spark job or stage.
Output Size / Records	Size of the file and number of records output by the Spark job or stage.
Status	<p>Status of the Spark job or stage. The status can be one of the following values:</p> <ul style="list-style-type: none"> - Running. The job or stage is running. - Success. The job or stage completed successfully. - Failed. The job or stage did not complete because it encountered errors. - Aborted. The job or stage did not complete because the user aborted the mapping task. <p>Note: After you abort a mapping task, there might be some lag time before the Monitor service shows the status as <i>Aborted</i>.</p>

Monitoring code tasks

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

Job properties

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

The following table describes the job properties for the code task:

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, code task.
CodeTask ID	Code task unique identifier.
Started By	Name of the user that started the job.

Property	Description
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.
Advanced configuration	Advanced configuration that was used to create the advanced cluster.
Cluster	Advanced cluster where the advanced 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 code task instance display the status of the job and success and error statistics.

The following table describes the job results for the code task:

Property	Description
Status	Job status. A job can have one of the following statuses: <ul style="list-style-type: none"> - Starting. The job is starting. - Running. The job is either queued or running. - Success. The job completed successfully. - Failed. The job did not complete because it encountered 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. Session log files are written to the following directory: <Secure Agent installation directory>/apps/Data_Integration_Server/logs
Requested Compute Units Per Hour	Number of serverless compute units per hour that the task requested. You can view the number of requested compute units if the task runs in a serverless runtime environment.
Total Consumed Compute Units	Total number of serverless compute units that the task consumed. You can view the number of consumed compute units if the task runs in a serverless runtime environment.
Error Message	Error message, if any, that is associated with the job.

Code task API execution parameters

The execution parameters for each code task instance display the API parameters used in the task.

The following table describes the execution parameters for the code task:

Property	Required / Optional	Description
Override Code Task Timeout	Optional	Overrides the code task timeout value for this execution. A value of -1 signifies no timeout.
Log Level	Optional	Log level for session logs, agent job log, Spark driver, and executor logs. Valid values are: none, terse, normal, verboseInitialization, or verboseData. The default value is normal.

The following table describes the Spark properties for the code task:

Property	Required / Optional	Description
Main Class	Required	Entry point of the Spark application. For example: org.apache.spark.examples.company.SparkExampleApp
Main Class Arguments	Optional	Ordered arguments sent to the Spark application main class. For example: --appTypeSPARK_PI_FILES_JARS-- classesToLoadcom.company.test.SparkTest1Class
Primary Resource	Required	Scala JAR file that contains the code task.
JAR File Path	Optional	The directory and file name of the JAR file that is uploaded to the cluster and added to the Spark driver and executor classpaths.
Spark File Path	Optional	The directory and file name of the Spark file that is uploaded to the cluster and available under the current working directory.
Custom Properties	Optional	Spark properties or other custom properties that Data Integration uses.

Spark application task details

The Spark application task details for each code task display under **Spark Application Task Results**.

Each Spark application task includes the following details:

Property	Description
Status	Status of the Spark task. The Spark task can have one of the following statuses: <ul style="list-style-type: none">- Running. The task is running.- Success. 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 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.
Duration	Amount of time that the Spark task ran.
Memory Per Executor	Amount of memory that each Spark executor uses.
Cores Per Executor	Number of cores that each Spark executor uses.
Driver and Agent Job Logs	Select Download to download the Spark driver and agent job logs.
Advanced Log Location	The log location that is configured in the advanced configuration for the advanced cluster. You can navigate to the advanced log location to view and download the agent job log, Spark driver log, and Spark executor logs.
Error Message	Error message, if any, that is associated with the job.

Each Spark application task is translated into Spark jobs, which are further broken down into stages. You can view the following details for each Spark job and stage:

Property	Description
Job Name	Name of the Spark job or stage.
Duration	Amount of time that the Spark job or stage ran.
Total Tasks	Number of tasks the Spark job or stage attempted.
Failed Tasks	Number of tasks that the Spark job or stage failed to complete.
Input Size / Records	Size of the file and number of records input by the Spark job or stage.

Property	Description
Output Size / Records	Size of the file and number of records output by the Spark job or stage.
Status	<p>Status of the Spark job or stage. The status can be one of the following values:</p> <ul style="list-style-type: none"> - Running. The job or stage is running. - Success. The job or stage completed successfully. - Failed. The job or stage did not complete because it encountered errors. - Aborted. The job or stage did not complete because the user aborted the code task. <p>Note: After you abort a code task, there might be some lag time before the Monitor service shows the status as <i>Aborted</i>.</p>

Monitoring dynamic mapping tasks

To view detailed information about a specific dynamic mapping task, click the instance name on the **My Jobs**, **All Jobs**, or **Running Jobs** pages.

When you view details for a completed task, you can restart the task run from the beginning. If one or more jobs in the task failed, you can resume the task run. When you resume a dynamic mapping task, Data Integration reruns jobs that failed or were skipped during the previous run. If the task run was stopped due to a warning and the task is configured to stop on warning, Data Integration also reruns jobs that completed with warnings.

You can also restart or resume a dynamic mapping task on the **My Jobs**, **All Jobs**, or **Running Jobs** pages.

Job properties

The job properties for each dynamic mapping task include general properties about the task.

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 fourth run of a task, this field displays "4."
Task type	Task type, for example, dynamic 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 that the job ran before it completed or stopped.
Runtime Environment	Runtime environment in which the job ran.

Job results

The job results for each dynamic mapping 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: <ul style="list-style-type: none">- 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..
Error Rows	Total number of source error rows, target error rows, and transformation errors.

Individual job results

The individual job results display results for jobs in the dynamic mapping task instance, and details for individual sources and targets.

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

Property	Description
Name	Name of the job, source transformation, or target object.
Group	Group the job is assigned to in the dynamic mapping task.
Start Time	Date and time that the job started.
End Time	Date and time that the job completed or stopped.
Success Rows	Number of rows successfully written to the target.
Affected Rows	Number of rows affected by the defined task operation.
Error Rows	Number of rows not written to the target.
Error Message	Error message, if any, that is associated with the job.
Status	Job status.

You can restart an individual job in a dynamic mapping task from the **My jobs**, **All jobs**, or **Running jobs** pages. To restart a specific job, in the row that contains the job, click **Restart**.

To see detailed information about each job, click the job name.

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: <ul style="list-style-type: none">- Starting. The job is starting.- Running. The job is still running.- Stopped. The job was stopped by a user or the parent job has stopped running, so the subtask can't start.- 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 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.
Requested Compute Units Per Hour	Number of serverless compute units per hour that the task requested. You can view the number of requested compute units if the task runs in a serverless runtime environment.
Total Consumed Compute Units	Total number of serverless compute units that the task consumed. You can view the number of consumed compute units if the task runs in a serverless runtime environment.
Error Message	Error message, if any, that is associated with the job.

Monitoring data transfer tasks

To view detailed information about a specific data transfer task, click the instance name on the **My Jobs**, **All Jobs**, or **Running Jobs** pages.

Job properties

The job properties for each data transfer task include general properties about the task.

The job properties include the following properties:

Properties	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, data transfer 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.

Properties	Description
Runtime Environment	Runtime environment in which the job ran.
Secure Agent	Secure Agent that ran the job.
Context Name	Name of the task that is associated with the job.

Job results

The job results for each data transfer 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: <ul style="list-style-type: none"> - 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.

Individual source and target results

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 object.
Success rows	Number of rows successfully written to the target.
Affected rows	Number of rows affected by the defined task operation.
Errors	Total number of source error rows or target error rows.
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.
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: <ul style="list-style-type: none">- Starting. The job is starting.- Running. The job is still running.- Success. The job completed successfully.- Stopped. The job was stopped by a user or the parent job has stopped running, so the subtask can't 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.

Property	Description
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
Requested Compute Units Per Hour	Number of serverless compute units per hour that the task requested. You can view the number of requested compute units if the task runs in a serverless runtime environment.
Total Consumed Compute Units	Total number of serverless compute units that the task consumed. You can view the number of consumed compute units if the task runs in a serverless runtime environment.
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.
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: <ul style="list-style-type: none">- 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 job was stopped by a user or the parent job has stopped running, so the subtask can't 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.
Requested Compute Units Per Hour	Number of serverless compute units per hour that the task requested. You can view the number of requested compute units if the task runs in a serverless runtime environment.

Property	Description
Total Consumed Compute Units	Total number of serverless compute units that the task consumed. You can view the number of consumed compute units if the task runs in a serverless runtime environment.
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."

Property	Description
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	<p>Job status. A job can have one of the following statuses:</p> <ul style="list-style-type: none"> - Starting. The job is starting. - Running. The job is still running. - Success. The job completed successfully. - Stopped. The job was stopped by a user or the parent job has stopped running, so the subtask can't 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.
Session Log	<p>Allows you to download the session log file. By default, Informatica Intelligent Cloud Services stores session logs for 10 runs before it overwrites the logs with the latest runs. If you need the session logs for earlier runs, take a backup of the directory that holds the session log files.</p> <p>Session log files are written to the following directory:</p> <pre><Secure Agent installation directory>/apps/Data_Integration_Server/logs</pre>
Requested Compute Units Per Hour	<p>Number of serverless compute units per hour that the task requested.</p> <p>You can view the number of requested compute units if the task runs in a serverless runtime environment.</p>
Total Consumed Compute Units	<p>Total number of serverless compute units that the task consumed.</p> <p>You can view the number of consumed compute units if the task runs in a serverless runtime environment.</p>
Error Message	Error message, if any, that is associated with the job.

Monitoring taskflows

You can monitor the details of the taskflows that are running or have run in your organization in the **All Jobs** and **Running Jobs** page. The **Started By** column displays the name of the user who first started the taskflow even if the taskflow is restarted by another user.


The taskflow instance name appears in the **All Jobs** and the **Running Jobs** page in the following format:

<taskflow name>-<runID>

If you had added a custom name to a taskflow name by using an API or the RunAJob utility, the taskflow name appears in the following format:

<taskflow name>-<custom name>-<runID>

The following image shows the **All Jobs** page that displays a taskflow instance with a custom name added to the taskflow name:



The screenshot shows the Informatica Monitor interface. The left sidebar contains navigation options: Running Jobs, All Jobs (selected), Import/Export Logs, File Transfer Logs, Source Control Logs, and Advanced Clusters. The main area displays the 'All Jobs' page with a table of taskflow instances. The table has columns for Instance Name, Location, Subtasks, Start Time, End Time, and Rows Processed. One instance, 'taskflowManagement-order-910067165634686976', is highlighted with a red box. Other instances include 'ACC_MCT-20225', 'tf_schedule_5min-910066682253725696', 'zdt_llhox_mct_dts-910066683084206080', and 'SimpleTask_Scheduler-910066682203394050'.

Instance Name	Location	Subtasks	Start Time	End Time	Rows Processed
taskflowManagement-order-910067165634686976	Harsha		Nov 17, 2023, 12:52...	Nov 17, 2023, ...	0
ACC_MCT-20225	Default		Nov 17, 2023, 12:50...	Nov 17, 2023, ...	0
tf_schedule_5min-910066682253725696	Default	1 task	Nov 17, 2023, 12:50...	Nov 17, 2023, ...	View Subtasks
zdt_llhox_mct_dts-910066683084206080	zdt_f	2 tasks	Nov 17, 2023, 12:50...	Nov 17, 2023, ...	View Subtasks
SimpleTask_Scheduler-910066682203394050	Hariya		Nov 17, 2023, 12:50...	Nov 17, 2023, ...	0

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

The following image shows the job details for a taskflow instance:

TaskflowParallelPCTasks-714788616458686464 | Success

1 Log

Step	Duration (HH:MM)
Start	00:00:0
Parallel Paths 1	00:00:2
Data Task 1	00:00:1
Data Task 2	00:00:2
Data Task 3	00:00:1
End	00:00:0

2 Design

3 Data Task 1 Properties | Success

4 Variable

Current Run

Status:	1 (Success)
Error:	No errors encountered.
Start Date:	2022/5/27 06:34:47.000 AM
End Date:	2022/5/27 06:34:58.000 AM
Duration (HH:MM:SS.ass):	00:00:11.000
Started By:	Taskflow: TaskflowParallelPCTasks
Username:	autouser_pod1

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.
4. Variable. The input and temp variables used to run the taskflow instance. To update the variables and their associated values, click the **Edit** icon next to the Input or Temp variable in the Variable panel.

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
Current Run	<p>You see the following properties:</p> <ul style="list-style-type: none">- Instance ID. A unique number that identifies the current taskflow instance.- Status. The state that the taskflow instance is in. Taskflows can be in the Running, Failed, Success or Suspended states.- Last Fault Name. The cause for the fault occurrence. This field appears only when the taskflow fails.- Error. Exception or error message.- Start Date. The date and time at which the taskflow instance was started.- End Date. 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.- 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.
Input Fields	Displays the name and value provided as input for the current taskflow instance.

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	<p>You see the following properties:</p> <ul style="list-style-type: none">- Status. Indicates the status 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 in the 0 (Faulted) or 2 (Warning) states, the error message appears here.- Start Date. The date and time that the data integration task started.- End Date. The date and time that 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 and type of the Data Task step.
Data Task	<p>You see the following properties:</p> <ul style="list-style-type: none">- Name. The name of the data integration task.- Type. Specifies whether the data integration task is a data transfer task, dynamic mapping task, mapping task, PowerCenter task, or synchronization task.

Property	Description
Input Fields	<p>Properties specific to the data integration task that you added to a Data Task step.</p> <p>You see the following properties:</p> <ul style="list-style-type: none"> - 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 (Seconds). The maximum length of time in seconds that the Data Task step waits for the data integration task to complete. - Task Name. The name of the data integration task. - GUID. The unique identifying number of the data integration task. - Task Type. Indicates whether the data integration task is a data transfer task, dynamic mapping task, mapping task, PowerCenter task, or synchronization task. - Has Inout Parameters. Indicates whether the data integration task contains input parameters or output parameters.
Output Fields	<p>You see output field properties if the State of the data integration task is 1 (Success).</p> <p>Properties specific to the completed data integration task appear here.</p> <p>Note: If the data task is a mapping task that is based on a mapping in advanced mode, the output field properties include success target rows and error target rows, but they do not include success source rows and error source rows.</p>
Fault	<p>If the data integration task faulted, you see the following properties:</p> <ul style="list-style-type: none"> - Code. Indicates whether the fault is an error or a warning. - Reason. Indicates the cause of the error or fault. - Detail. Provides details about the error or fault. <p>You see the output field properties when the data integration task status is 3 (Failed) and one of the following conditions are met:</p> <ul style="list-style-type: none"> - The On Error field is set to Ignore or Custom error handling. - The Fail taskflow on completion option is set to If this task fails.

IntegrationOps Task properties

The IntegrationOps Task step properties display details of the current run of the IntegrationOps task in a taskflow. The taskflow continues to run the subsequent steps even if the IntegrationOps task fails.

When you click the IntegrationOps task, the **Properties** panel displays the following properties:

Property	Description
Input Fields	<p>Properties specific to the process you added to the IntegrationOps Task step.</p> <p>Displays the following properties:</p> <ul style="list-style-type: none"> - Task Name. The name of the Application Integration process. - GUID. The unique identifying number of the process. - Task Type. The task type. For an IntegrationOps Task, the task type is always PROCESS. - Agent Name. The Secure Agent name where the process ran.
Output Fields	<p>Displays the output fields of the process along with their associated values.</p> <p>The output fields are displayed only if the IntegrationOps task completed successfully.</p>
Fault	<p>If the IntegrationOps task is faulted, you see the following properties:</p> <ul style="list-style-type: none"> - Code. Indicates whether the fault is an error or a warning. - Reason. Indicates the cause of the error or fault. - Detail. Provides 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	<p>Displays the following input fields of the Notification Task:</p> <ul style="list-style-type: none">- 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 Bcc. The additional recipients who need to be sent a copy of the email notification. The recipients in the Email To and Email Cc fields will not be able to see the recipients in the Email Bcc field. If the field contains more than one recipient, the recipients will not be able to see the other recipients in the Email Bcc field.- 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. <p>If you selected the email body as Content, the email content appears within HTML tags. For example, <code><p>Hello World</p></code></p>

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. The fields also display the maximum length of time in seconds that the command task step waits for the task to complete.
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 value of the run ID, start and end time, exit code, execution status, and standard error. The fault fields are displayed only if the command task fails due to a script failure. The details help you analyze the reason for the fault occurrence. You can then take an appropriate action on the faulted command task and proceed with the execution of the taskflow.

If the Command Task step contains multiple script files, you can view all the scripts in the taskflow instance as subtasks of the command task from the **All Jobs**, **Running Jobs**, or **My Jobs** page. If a script fails, the status of the script is set to failed. You can also download the log file to understand the reason for the script failure.

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: <ul style="list-style-type: none">- 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 .
Input Fields	Displays the following input fields: <ul style="list-style-type: none">- 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/.

Property	Description
Output Fields	<p>Displays the following output fields if the File Watch Task step completed successfully:</p> <ul style="list-style-type: none"> - 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: <ul style="list-style-type: none"> - Event ID. Auto-generated ID for the event. - Event type. Indicates the nature of the file event. You see one of the following values: <ul style="list-style-type: none"> - 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	<p>If the file watch task faulted, you see the following properties:</p> <ul style="list-style-type: none"> - 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	<p>Displays the following details:</p> <ul style="list-style-type: none"> - Status. Status of the file ingestion task. - Start Date. Date and time when the file ingestion task started. - End Date. Date and time when the file ingestion task ended or was suspended. - Duration. Amount of time in seconds that the file ingestion task ran before it completed or was suspended. - Started By. Name of the taskflow that called the file ingestion task.
General	Displays the name of the Ingestion Task step.
Ingestion Task	<p>Displays the name and type of the ingestion task that the Ingestion Task step contains.</p> <p>The type is set to MI_TASK for file ingestion tasks.</p>
Input Fields	<p>Displays the following input fields:</p> <ul style="list-style-type: none"> - Wait for Task to Complete. Indicates whether the Ingestion Task step waits for the file 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 file ingestion task to complete. This value is always set to 604800 seconds, which is 7 days. - Task Name. Name of the file ingestion task. - GUID. Unique identification number for the file ingestion task. - Task Type. Indicates the type of task that the Ingestion Task step contains. This value is always set to MI_TASK for file ingestion tasks. - serviceProviderId. Service provider ID for the file ingestion task. This value is always set to filelistener. - serviceProviderContextPath. Service provider context path for the file ingestion task. This value is always set to mitask/.
Output Fields	<p>Displays the following output fields if the Ingestion Task step completed successfully:</p> <ul style="list-style-type: none"> - endTime. Date and time when the file ingestion task completed. - failedFiles. Number of failed files in the file ingestion task run. - jobNumber. ID for the file 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 file ingestion task started. - status. Indicates whether the Ingestion Task step completed. This value is always set to Success. - successFiles. Number of files that were successfully processed in the file ingestion task. - fileDetails. Provides a list of objects with the following details: <ul style="list-style-type: none"> - Duration. Amount of time in seconds that the file 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: <ul style="list-style-type: none"> - DOWNLOAD. The file was downloaded from the source. - UPLOAD. The file was uploaded to the target.
Fault	<p>If the ingestion task faulted, you see the following properties:</p> <ul style="list-style-type: none"> - 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.

Subtaskflow step properties

The Subtaskflow step properties display details about the current run of the subtaskflow.

When you click a Subtaskflow step, the **Properties** panel displays the following information:

Properties	Description
Input Fields	Displays the input fields that you had configured along with their associated values.
Fault	If the subtaskflow step faulted, you see the following properties: <ul style="list-style-type: none">- Code. Indicates whether the fault is an error or a warning.- Reason. Indicates the cause of the error or fault.- Source. Provides details about the error or fault. If a Data Task step that uses a mapping task fails, the subtaskflow name followed by the mapping task name is displayed.
Output Fields	Displays the output fields of the subtaskflow along with their associated values.

Throw step properties

The Throw step properties display the fault details for the taskflow or the step that is associated with the Throw step.

If the step preceding the Throw step or the step with which the Throw step is associated faulted, the **Fault** field in the **Properties** panel displays the following details:

- **Code**. Displays the fault code.
- **Reason**. Displays the cause of the fault.
- **Details**. Displays the details about the fault.

End step properties

The End step properties display details about the current run of the taskflow.


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

Properties	Description
Fault	If the End step faulted, you see the following properties: <ul style="list-style-type: none">- Code. Displays the fault code. The code is set to faultDueToNotRunOrFaulted if the subtaskflow associated with the Subtaskflow step failed or did not run.- Reason. Indicates the cause of the error or fault.- Source. Displays the path of the step that faulted

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.
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: <ul style="list-style-type: none">- Starting. The job is starting.- Running. The job is still running.- Success. The job completed successfully.- Stopped. The job was stopped by a user or the parent job has stopped running, so the subtask can't 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	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: <ul style="list-style-type: none">- 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 subtask was stopped by a user or the parent job has stopped running, so the subtask cannot start. If you stop a subtask, Data Integration does not propagate the status to the parent task. <ul style="list-style-type: none">- Warning. The job completed with errors.- Failed. The job did not complete because it encountered errors.- Aborted. The job was aborted. Applies to file ingestion and replication subtasks.
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.

Property	Description
Error message	Error encountered when writing to the target, if any.
Recommendation	Recommended action for errors.

Running only a particular step in a taskflow instance

You can run only a particular step in a taskflow instance from the **All Jobs**, **Running Jobs**, or **My Jobs** page.

When you run a particular step in a taskflow instance, Data Integration locks the entire taskflow so that other users can't run the taskflow. Data Integration runs only the selected step.

You can run only a particular step if the taskflow meets the following criteria:

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


You can run a taskflow instance from any step other than the Start step, End step, and Jump step.

You can run only the following steps in a taskflow instance as a standalone step:


- Data Task step
- Ingestion Task step
- Command Task step
- IntegrationOps step
- Notification Task step
- File Watch Task step
- Subtaskflow step

To run a particular step in a taskflow instance, open the **All Jobs**, **Running Jobs**, or **My Jobs** page, and then click the taskflow instance. On the job details page, select the step that you want to run and click the **Run only the selected step** icon.

The following image shows the **Run only the selected step** icon to run the selected Notification Task step in the taskflow instance:



The following image shows the taskflow instance after it runs the selected Notification Task step:



When you run a failed task, if the task fails again at rerun, the overall status of the taskflow instance appears as failed on the job details page and on the **All Jobs**, **Running Jobs**, and **My Jobs** pages.

Note: The taskflow ignores the error handling setting for the failed task if set to ignore on error.

Running a taskflow instance from a step

You can run a taskflow instance from a particular step from the **All Jobs**, **Running Jobs**, or **My Jobs** page.

You can run a taskflow instance from a step if the taskflow meets the following criteria:

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

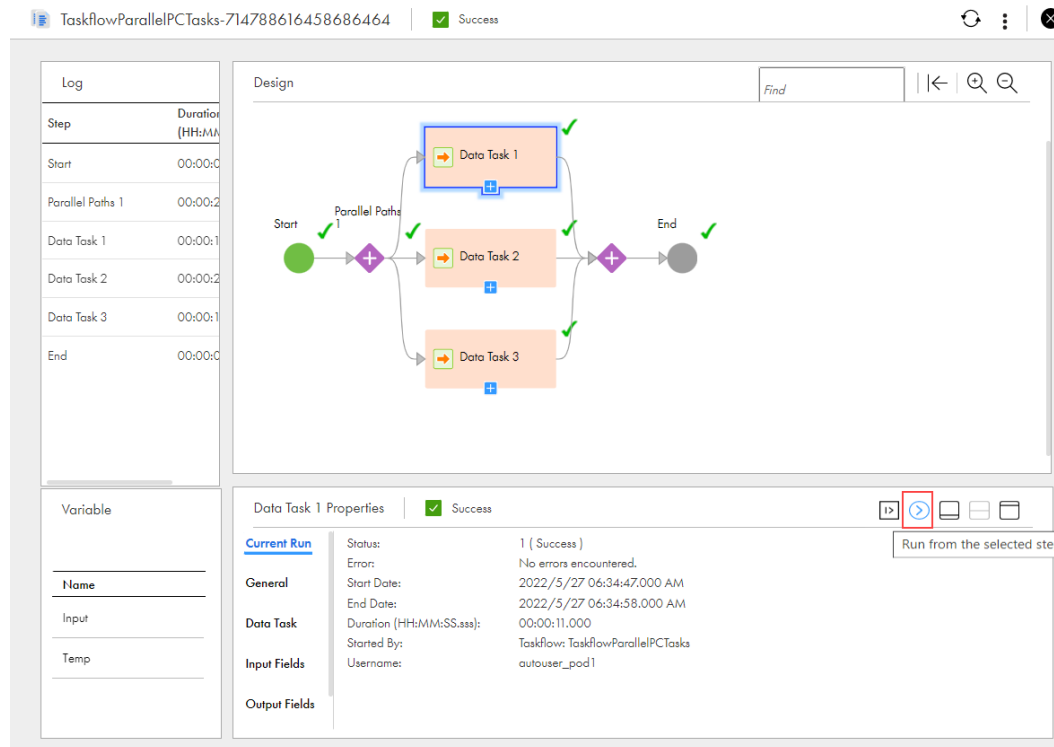
You can run a taskflow instance from any step other than the Start step, End step, Throw step, and Jump step.

When you run a taskflow instance from a particular step, if any subsequent step uses input from a previously executed step, it will have null values. To avoid this issue, update the input and temporary variables before

you run a taskflow instance from a step. For more information about updating variables, see [“Updating variables to run a taskflow instance from a step” on page 58](#).

To run a taskflow instance from a step, open the **All Jobs**, **Running Jobs**, or **My Jobs** page, and then click the taskflow instance. On the job details page, select the step from which you want to run the taskflow and click the **Run from the selected step** icon.

The following image shows the **Run from the selected step** icon to run the taskflow instance from the selected Data Task step:



Note: When you run a taskflow from a particular step, the taskflow status might change. If the taskflow is used in another taskflow, the parent taskflow will not be impacted by the updated taskflow status. You can run a taskflow from the steps that ran successfully. You can also run a taskflow from a failed task such as a Data Task step, Ingestion Task step, or Command Task step provided the error handling for that task is set to ignore on error.


Updating variables to run a taskflow instance from a step

Use the **Variables** tab to update the input and temporary variables or fields and their associated values to run a complex taskflow instance from a particular step.

The **Variables** tab contains all the input variables and fields that are configured in the Start step for the taskflow. It also contains the temporary variables and fields that also include tasks such as Data Task step, Ingestion Task step, and Command Task step that were used to run the taskflow.


To update the input and temporary variables or fields and their associated values, open the **All Jobs** and **Running Jobs** page in Monitor or the **My Jobs** page in Data Integration, and then click the taskflow instance. On the job details page, click the **Edit** icon next to the Input or Temp variable. Based on the values passed in the taskflow, the variables are displayed in the JSON or XML format in the **Variable Input** and **Variable Temp** dialog boxes respectively. Update the variables and click **Update**.

You cannot run a task that never ran. For example, you cannot run a taskflow instance from the Data Task 1 step as shown in the following image:



In this case, you must update the variable value that is used in the Decision 1 step, and then run from the Decision 1 step again.

The following image shows the taskflow instance after it runs from the Decision 1 step:



For more information about the input and temporary fields, see the Taskflows documentation.

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.

Stop the job using one of the following methods:

- **Clean Stop.** Gracefully shuts down the DTM process.
- **Stop.** Forcefully stops the DTM process.

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

When you click **Clean Stop**, the job status changes to Stopping. If the job takes too long to stop, click **Stop** to stop the job.

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.

You can also stop the taskflow instance from the job details page. To stop a taskflow instance, open the **All Jobs**, **Running Jobs**, or **My Jobs** page, and then click the taskflow instance. On the job details page, click the **Stop** icon at the top of the page.

When you stop a parent taskflow, all the underlying subtaskflows that are running also stop. The **Last Fault Name** field in the **Taskflow Properties > Current Run** section of the parent taskflow that you stop displays the value as **processManuallyTerminated**. However, the **Last Fault Name** field of the subtaskflows displays the value as **processTerminated**.

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 from a faulted step.

To resume a suspended taskflow instance from a faulted step, open the **All Jobs**, **Running Jobs**, or **My Jobs** page, and then click the taskflow instance. On the job details page, click the **Resume from faulted step** icon. The taskflow resumes from the step that faulted.

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.

Skip a faulted step and resume a taskflow instance from the next step.

To skip a faulted step and resume a suspended taskflow instance from the next step, open the **All Jobs**, **Running Jobs**, or **My Jobs** page, and then click the taskflow instance. On the job details page, click the **Skip faulted step and resume** icon. The taskflow skips the faulted step and resumes from the next step. If all the other steps run successfully, the taskflow status is set to success.

Restart a taskflow instance.

Users who have the Admin, Designer, 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.

Taskflow log files

You can download the log file for a taskflow to review it offline and troubleshoot problems, if necessary. The downloaded log file name contains the name of the taskflow, the run ID, and the organisation ID of the user who triggered the taskflow.

Note: If the taskflow name contains a space, the log file name contains a + symbol.

Downloading a taskflow log file from Monitor

You can download a taskflow log file from Monitor to review it offline and troubleshoot problems.

1. In Monitor, open the **All Jobs** page or the **Running Jobs** page.
2. Click the taskflow in the **Instance Name** column to open the corresponding taskflow details page.
3. In the **Results** area, click the **Download Log** icon to download the taskflow log file.

Downloading a taskflow log file using the log resource

After you publish a taskflow as a service, you can run the taskflow and download the log file for the taskflow log file using the log resource.

You can use one of the following ways to invoke the request through a REST client such as Postman:

- Use basic authentication, and enter the user name and password.
- Use the session ID in the HTTP header to invoke the taskflow without providing the user name and password.

For example, you can authenticate the GET request in one of the following ways:

- Use basic authorization and specify the Informatica Intelligent Cloud Services user name and password.
For example:

```
GET<Informatica Intelligent Cloud Services URL>/active-bpel/services/tf/log/<run ID>  
Accept: application/octet-stream  
Authorization : Basic Auth
```

```
username: <Informatica Intelligent Cloud Services username>
password: <<Informatica Intelligent Cloud Services password>
```

- Use the IDS-SESSION-ID in the HTTP header.

For example:

```
GET<Informatica Intelligent Cloud Services URL>/active-bpel/services/tf/log/<run ID>
Accept: application/octet-stream
IDS-SESSION-ID : <sessionId>
```

To get the SESSION-ID, use the Platform REST API version 3 login resource. For more information about the login resource, see *REST API Reference*.

GET request

To download a log file for the taskflow, use the following URI:

```
GET<Informatica Intelligent Cloud Services URL>/active-bpel/services/tf/log/<run ID>
```

For example:

```
https://na4.dm.us.informaticacloud.com/active-bpel/services/tf/log/20262247166322413568
```

Taskflow log file contents

The header of the log file contains information such as the time stamp when the file was downloaded, user ID, and the status of the taskflow at the time of downloading the log file. The body of the log file contains the following information:

Field	Type	Description
assetName	String	Name of the taskflow. The taskflow name also includes the custom name, if you had added a custom name to the taskflow using an API or the RunAJob utility.
assetType	String	Type of the object. Returns the value TASKFLOW.
duration	String	Time in seconds that the taskflow ran before it completed, was suspended, was failed, or was stopped.
endTime	Date/time	End time for the taskflow run. Uses Coordinated Universal Time (UTC).
location	String	Project and folder path where the taskflow is saved.
runId	Long	Run ID for the taskflow.
URL	String	The log resource URL generated by Data Integration. The log file will contain one URL corresponding to the taskflow and as many URLs as the subtasks within the taskflow.
runtimeEnv	String	The name of the application node where the taskflow runs.
runtimeEnvName	String	Name of the runtime environment where the taskflow runs.
startTime	Date/time	Start time for the taskflow run. Uses Coordinated Universal Time (UTC).
startedBy	String	Name of the user who first started the taskflow.

Field	Type	Description
status	String	Execution status of the taskflow. Returns one of the following values: - RUNNING. The taskflow is running. - SUCCESS. The taskflow completed successfully. - FAILED. The taskflow did not complete because it encountered errors. - SUSPENDED. The taskflow run was suspended.
subtasks	String	Number of subtasks that the taskflow contains.
updateTime	Date/time	Last time the taskflow run status was updated. Uses Coordinated Universal Time (UTC).
statusCode	String	Status code of the taskflow. The status code can be one of the following values: - 1 - Running or CHILD_SUSPENDED - 2 - Suspended - 3 - Success - 4 - Failed - -1 - Default
errorMessage	String	Error message string.
subtaskDetails	String	Object that contains status details for all subtasks in the taskflow.
details	String	Status details. Includes status information for each subtask in the tasks object.
tasks	Collection	Status information for all subtasks that the taskflow contains.

Sample taskflow log file

The content of the log file is of octet-stream type. The following snippet shows a sample of the downloaded log file for a taskflow:

```
{
  [Tue Aug 29 09:49:11 UTC 2023] [sadityaRelPod1] - TASKFLOW
  TaskflowRelease15467983 FAILED.
  Details:
  {
    "assetName" : "TaskflowRelease15467983",
    "assetType" : "TASKFLOW",
    "duration" : 795,
    "endTime" : "2023-08-29T06:05:44Z",
    "errorMessage" : "{\"reason\":\"Unknown\",\"code\":\"Unknown\",\"details\":\"Unknown\"}",
    "location" : "Default", "runId" : 881053605371789312,
    "URL" : "https://qa-pod1.rel.infaqa.com/active-bpel/services/tf/log/881053605371789312",
    "runtimeEnv" : "taskflow-qa-release-pod1-r41-app02.infacloudops.net:4430",
    "runtimeEnvName" : "",
    "startedBy" : "sadityaRelPod1",
    "startTime" : "2023-08-29T05:52:29Z",
    "status" : "FAILED",
    "subtasks" : 1,
    "updateTime" : "2023-08-29T06:05:44Z",
    "statusCode" : 4,
    "subtaskDetails" : {
      "details" : {
        "tasks" : [ {
          "assetName" : "MappingTask2",
          "assetType" : "MTT",
          "duration" : 13,
          "endTime" : "2023-08-29T05:52:43Z",
```

```
"errorMessage" : "",
"errorRows" : 0,
"location" : "Shivendra",
"rowsProcessed" : 6,
"runId" : 5,
"runtimeEnv" : "01589925000000000002",
"runtimeEnvName" : "INWPG02P31F-AAD",
"startedBy" : "sacityarelpod1",
"startTime" : "2023-08-29T05:52:30Z",
"status" : "COMPLETED",
"subtasks" : 0,
"successRows" : 6,
"updateTime" : "2023-08-29T05:52:44Z"    } ]    } }}
}
```


CHAPTER 3

Data Integration job log files

Data Integration generates log files to help you monitor running, failed, and completed jobs. You can access some of the log files from the **All Jobs**, **Running Jobs**, and **My Jobs** pages, and 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. This log gives you a high level view of the time spent for different operations.

The session log provides mapping compilation time, translation time, simplification time, optimization time, total time to create the LDTM, Spark task submission time, Spark task [InfaSpark0] execution start and end time, and Total time to perform the LDTM operation.

If a job fails, analyze the session log file first to help you troubleshoot the job.

Reject file

Data Integration creates a reject file for each flat file and Oracle target in a mapping or mapping task that contains error rows. The reject file contains information about each rejected target row and the reason that the row was rejected. Data Integration saves the reject file to the following default folder:

```
$PMBadFileDir/<task federated ID>
```

Execution plan

Data Integration generates an execution plan that shows the Scala code that an advanced cluster uses to run the data logic in a mapping in advanced mode. You can use the Scala code to debug issues in the mapping.

Agent job log

Data Integration generates an agent job log that shows the logic that the Secure Agent uses to push the Spark execution workflow for a mapping in advanced mode to an advanced cluster for processing.

The agent job log contains information such as metering, time the application was submitted to the cluster, and the time the application completed. This log can help you troubleshoot delays in running the Spark task in the Session log, and you can see when the Spark task was processed on the Secure Agent.

Spark driver and Spark executor logs

An advanced cluster generates Spark driver and Spark executor logs to show the logic that the cluster uses to run a job. Use these logs to identify issues or errors with the Spark process. This log also contains information about the different executors being created and the tasks that are starting or have been completed.

Initialization script log

If an initialization script runs on an advanced cluster, the cluster generates an init script log to show the script output.

Cloud-init log

If an initialization script runs on the advanced cluster, the cluster generates a cloud-init log that 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.

Spark event log

An advanced cluster generates a Spark event log to stream runtime events for tasks that run on the cluster.

The Spark event log records different events in a JSON-encoded format while the application is running. This log contains the events associated with the Spark application, such as the different jobs in the application, different stages, individual tasks, and interaction between entities.

The Spark event log also contains events related to the software infrastructure like driver information, executor creation, memory usage by executors, environment configuration, and the logical and physical plans of the Spark application. Use this log to trace what happened during every step of the Spark application run.

To find the Spark event log, open the Spark driver log and search for `SingleEventLogFileWriter`. The result of the search shows the path of the Spark event log. For example:

```
23/01/09 04:38:35 INFO SingleEventLogFileWriter - Logging events to s3://bucket/
log_location_in_cluster_condifuration/eventLogs/atyscaleagent/spark-
a7bea557ede14382b4807d35b5404b97.inprogress
```

When the application completes, download the Spark event log from the location `s3://bucket/log_location_in_cluster_condifuration/eventLogs/atyscaleagent/` as file `spark-a7bea557ede14382b4807d35b5404b97`.

To interpret the Spark event log, import it into a Spark history server and examine the log using the History server monitor. Check the following tabs:

- The **Jobs** tab shows all the detailed metrics.
- The **Stages** tab lists all the completed stages. You can see detailed information on the total number of tasks succeeded or failed, input and output data volume, and shuffle read and shuffle write data volume. Click on any stage to view the DAG visualization diagram.
- The **Environments** tab shows the Spark-related parameters used to run the Spark job.
- The **Executors** tab shows detailed information about the executor Pods and driver Pod.

For more information, refer to the Apache Spark documentation.

Advanced logs

The Advanced Log Location contains the Spark executor logs apart from the Spark driver and Agent job logs. The executor logs can help you troubleshoot issues with individual executors.

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 advanced cluster subtasks

When you monitor a subtask that runs on an advanced cluster, 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 the 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
Spark event 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 in the Spark task details in Monitor. You can also 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, you can download the advanced logs from Monitor. The **Download** option becomes available when the job completes.

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 advanced configuration.

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

Downloading Spark event logs

To download the Spark event log, navigate to the following location:

```
<Log location>/eventLogs/atscaleagent
```

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

If the job fails to start, you see the following sample text in the event log file:

```
{"Event":"SparkListenerLogStart","Spark Version":"3.0.0"}
```

For more information about the job failure, see the Spark driver log.

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 a subtask that runs on an advanced cluster 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 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	<p>Job status. A job can have one of the following states:</p> <ul style="list-style-type: none"> - 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. <p>When a job is in Running status, a Stop icon appears in the Status property. Click the icon to stop the job.</p>

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 <i>Data Profiling</i> . 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.

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

- Fetching the source row count
- s_profiling
- Drilldown
- Query

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

You can view the following details for the Secure Agent in the session log file for the profile mapping jobs:

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

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 5

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.

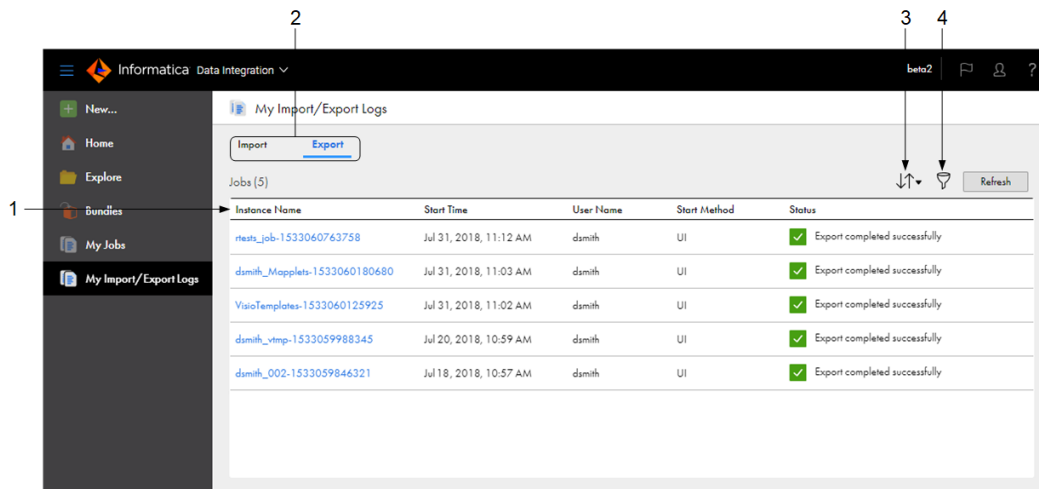
Property	Description
User name	Name of the user that started the import or export.
Start method	Method in which the import or export was started. An import or export instance can have one of the following start methods: <ul style="list-style-type: none"> - UI. The import or export was started from the 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 screenshot shows a window titled "Export_mt_FilterLACustomers-1511828517829". The main content area is divided into two sections: "Properties" and "Exported Assets (2)".

Properties:

Name:	Export_mt_FilterLACustomers-1511828517829	Started By:	ltroy
Type:	Export	Start Method:	UI
Start Time:	Jul 20, 2018, 10:04 AM	Status:	✔ Export completed successfully
End Time:	Jul 20, 2018, 10:05 AM	Export File:	Export_mt_FilterLACustomers-1511828517829.zip
		Export Log:	Download Export Log

Exported Assets (2):

Name	Type	Source Location	Description	Status
m_FilterLACustomers	Mapping	Development\Customers	Exclude customer records in whic...	✔
mt_FilterLACustomers	Mapping Task	Development\Customers	Run mapping m_FilterLACustome...	✔

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: <ul style="list-style-type: none">- 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 6

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 screenshot shows a table titled "Jobs (14)" with columns for Instance Name, Asset Type, Command, End Time, # of Files, Sender/Receiver ID, MDN, and Status. The table contains 14 rows of log entries, most with a "Success" status and a green checkmark icon, and one with a "Failed" status and a red X icon.

Instance Name	Asset Type	Command	End Time	# of Files	Sender/Receiver ID	MDN	Status
as2_automation-1002	AS2 Receive	Transfer	Jun 7, 2019, 6:44 AM	1	as2_automation	Sent	Success
as2_automation-1001	AS2 Receive	Login	Jun 7, 2019, 6:44 AM	0			Success
AS2-134	AS2 Send	Transfer	Jun 9, 2019, 11:08 PM	2	as2_automation	Email	Success
AS2-133	AS2 Send	Transfer	Jun 9, 2019, 11:05 PM	2	as2_automation	Received	Success
AS2-132	AS2 Send	Transfer	Jun 9, 2019, 11:03 PM	0	as2_automation		Success
AS2-131	AS2 Send	Transfer	Jun 9, 2019, 10:59 PM	2	as2_automation	Received	Success
AS2-127	AS2 Send	Transfer	Jun 7, 2019, 6:44 AM	1	as2_automation	Received	Success
AS2-126	AS2 Send	Transfer	Jun 7, 2019, 6:31 AM	0	as2_automation		Failed
as2-1008	AS2 Receive	Transfer	Jun 9, 2019, 11:08 PM	2	as2	Sent	Success
as2-1007	AS2 Receive	Login	Jun 9, 2019, 11:08 PM	0			Success
as2-1006	AS2 Receive	Transfer	Jun 9, 2019, 11:05 PM	2	as2	Sent	Success
as2-1005	AS2 Receive	Login	Jun 9, 2019, 11:05 PM	0			Success
as2-1004	AS2 Receive	Transfer	Jun 9, 2019, 10:59 PM	2	as2	Sent	Success
as2-1003	AS2 Receive	Login	Jun 9, 2019, 10:59 PM	0			Success

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: <ul style="list-style-type: none">- 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: <ul style="list-style-type: none">- 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:

Instance Name	Asset Type	Command	End Time	# of Files	Status
SFTP DD1-1035	SFTP Server	Login	Oct 4, 2019, 5:25...	0	Success
SFTP DD1-1036	SFTP Server	Disconnect	Oct 4, 2019, 5:30...	0	Success
SFTP DD1-1038	SFTP Server	Login	Oct 4, 2019, 5:30...	0	Success
SFTP DD1-1039	SFTP Server	Disconnect	Oct 4, 2019, 5:35...	0	Success
SFTP DD1-1041	SFTP Server	Login	Oct 4, 2019, 5:35...	0	Success
SFTP DD1-1042	SFTP Server	Disconnect	Oct 4, 2019, 5:40...	0	Success

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:

Instance Name	Asset Type	Command	End Time	# of Files	Status
https_automation-6370	HTTPS Server	Login	Jan 29, 2020, 5:10 PM	0	Success
https_automation-6368	HTTPS Server	Login	Jan 29, 2020, 5:10 PM	0	Success
https_automation-6369	HTTPS Server	Login	Jan 29, 2020, 5:10 PM	0	Success
https_automation-6367	HTTPS Server	Logout	Jan 29, 2020, 5:10 PM	0	Success
https_automation-6364	HTTPS Server	Login	Jan 29, 2020, 5:10 PM	0	Success
https_automation-6365	HTTPS Server	Upload	Jan 29, 2020, 5:10 PM	1	Success
https_automation-6366	HTTPS Server	Upload	Jan 29, 2020, 5:10 PM	1	Success
https_automation-6363	HTTPS Server	Logout	Jan 29, 2020, 5:10 PM	0	Success
https_automation-6360	HTTPS Server	Login	Jan 29, 2020, 5:10 PM	0	Success
https_automation-6361	HTTPS Server	Upload	Jan 29, 2020, 5:10 PM	1	Success
https_automation-6362	HTTPS Server	Upload	Jan 29, 2020, 5:10 PM	1	Success
https_automation-6359	HTTPS Server	Logout	Jan 29, 2020, 5:10 PM	0	Success
https_automation-6357	HTTPS Server	Upload	Jan 29, 2020, 5:10 PM	1	Success
https_automation-6358	HTTPS Server	Upload	Jan 29, 2020, 5:10 PM	1	Success
https_automation-6356	HTTPS Server	Login	Jan 29, 2020, 5:09 PM	0	Success
ais2_cloud-6355	HTTPS Server	Login	Jan 29, 2020, 5:07 PM	0	Success
https_automation-6327	HTTPS Server	Logout	Jan 29, 2020, 11:28 AM	0	Success


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 MLLP file transfers

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

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

File Transfer Logs page

Lists all file transfers that are in progress or have completed. You can filter the list to view only the MLLP transfers.

Details page

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

Viewing MLLP file transfer logs

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

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

The **File Transfer Logs** page lists MLLP file transfer logs and provides status information. It also displays the acknowledgement type.

The following image shows the **File Transfer Logs** page filtered for MLLP logs:

Instance Name	Asset Type	ACK Type	Message ID	Event Type	End Time	Status
130750	MLLP Server	AE	YhMUBGe4WBuGh2...	MLLP Message Receiv...	Feb 1, 2023, 1...	Failed
130749	MLLP Server	AA	YhMUBGe4WBuGh2...	MLLP Message Receiv...	Feb 1, 2023, 1...	Success
130252	MLLP Server	AA	YhMUBGe4WBuGh2...	MLLP Message Receiv...	Jan 31, 2023, ...	Success
129027	MLLP Server	AA	0123456789	MLLP Message Receiv...	Jan 30, 2023, ...	Success
128993	MLLP Server	AA	id201	MLLP Message Receiv...	Jan 30, 2023, ...	Success
128391	MLLP Server	AA	0123456789	MLLP Message Receiv...	Jan 30, 2023, ...	Success
128390	MLLP Server	AA	0123456789	MLLP Message Receiv...	Jan 30, 2023, ...	Success
127235	MLLP Server	AA	3201	MLLP Message Receiv...	Jan 26, 2023, ...	Success
127234	MLLP Server	AA	3201	MLLP Message Receiv...	Jan 26, 2023, ...	Success
127233	MLLP Server	AA	3201	MLLP Message Receiv...	Jan 26, 2023, ...	Success
127232	MLLP Server	AA	3201	MLLP Message Receiv...	Jan 26, 2023, ...	Success


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

Property	Description
Instance Name	Name of the MLLP file transfer instance.
Asset Type	MLLP Server.
ACK Type	Acknowledgement type for the log. The ACK type can be one of the following: - AA: Application accept - AE: Application Error - AR: Application Reject
Message ID	Identifier for the message.
Event Type	Type of MLLP event.
End Time	Date and time that the action executed by the command ended.
Status	Status of the action, success or failed.

Viewing details of an MLLP file transfer log

To view detailed information about an MLLP file transfer log, click the instance name on the **File Transfer Logs** page. When you click the instance name, details about the MLLP event appears.

The following image shows the details of an MLLP event:



The screenshot shows a window titled "MLLP 130749" with a "Refresh" button and a close icon. The main content is titled "Event Properties" and lists the following details:

- Server IP: 0.0.0.0
- Server Port: 2579
- Remote IP: 10.65.41.187
- Remote Port: 57349
- Start Time: Feb 1, 2023, 12:37:40 PM
- End Time: Feb 1, 2023, 12:37:40 PM
- Duration: 00:00:01
- Event Type: MLLP Message Receive Successful
- Message ID: YhMUBGe4WBUrGh2LdVWxCr4jrp
- ACK Type: AA
- Target Location: Kafka Topic : MREL-Test-MLLP
- State: Success
- Remarks: MLLP Message received successfully


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:



The screenshot shows a window titled "ftp_to_ftp_file_1 FriSep042020145255-1092" with a "Refresh" button. The main content is titled "File Listener Job Properties" and is divided into two sections: "File Listener Job Properties" and "Results".

File Listener Job Properties:

- Asset Name: ftp_to_ftp_file_1 FriSep042020145255
- Asset Type: File Listener
- Started By: b2b
- Start Time: Sep 4, 2020, 2:34:09 AM
- End Time:
- App Version: 11.15.1
- Snapshot name: false

Results:

- Status: Running
- Arrived Files: 7
- Updated Files: 0
- Deleted Files: 0
- Remarks: File Listener has been started by user: b2b

File Events (7):

Name	File Size (Bytes)	Last Modified	Path	Event Type	Event Time
File1.txt	47	Dec 23, 2019, 3:58:59 AM	/root/.suricata/logs/automation/KSFfile/File1.txt	Arrives	Sep 4, 2020, 2:24:58 AM
File2.txt	47	Jul 5, 2020, 10:15:37 PM	/root/.suricata/logs/automation/KSFfile/File2.txt	Arrives	Sep 4, 2020, 2:24:58 AM
File3.txt	47	Dec 23, 2019, 3:58:57 AM	/root/.suricata/logs/automation/KSFfile/File3.txt	Arrives	Sep 4, 2020, 2:24:58 AM
File4.txt	28	Dec 23, 2019, 3:58:57 AM	/root/.suricata/logs/automation/KSFfile/File4.txt	Arrives	Sep 4, 2020, 2:24:58 AM
File5.txt	37	Dec 23, 2019, 3:58:57 AM	/root/.suricata/logs/automation/KSFfile/File5.txt	Arrives	Sep 4, 2020, 2:24:58 AM
File6.txt	35	Dec 23, 2019, 3:58:57 AM	/root/.suricata/logs/automation/KSFfile/File6.txt	Arrives	Sep 4, 2020, 2:24:58 AM
File7.txt	49	Dec 23, 2019, 3:34:09 AM	/root/.suricata/logs/automation/KSFfile/File7.txt	Arrives	Sep 4, 2020, 2:24:58 AM

File listener job details

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.

Job properties

The job properties for each file listener job instance display general properties about the instance.

The job properties include the following properties:

Status	Reason
Asset Name	Name of the asset.
Asset Type	The type of asset. In this case, file listener.
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.
App Version	Agent version on which the file listener instance is running.
Snapshot reset	When you update the file listener definition, such as updating the run-time environment, connector type, connection, or listener rule properties, the file listener requires a reset. This action clears the list of previously tracked files. If the file listener is reset, the snapshot reset value is true.

Job results

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

The job results include the following properties:

Status	Reason
State	A job can have one of the following statuses: <ul style="list-style-type: none">- 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:<ul style="list-style-type: none">- 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 with Error: The file listener has encountered errors.- Running: The file listener run is in progress.
Arrived Files	Number of files that arrived to the target.
Updated Files	Number of files that were updated.

Status	Reason
Deleted Files	Number of files that were deleted.
Remarks	This is auto populated, and it is based on the state of the job. For example, if the file listener is in Failed status, the message includes the reason for failure.

Note: The file listener fails if the number of file events such as arrive, update, or delete, exceeds the maximum event limit.

File events

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

The File Events section is updated each time the file listener job 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:

Status	Reason
Name	The name of the file.
File Size (Bytes)	The size of the file in bytes.
Last Modified	The date and time when the file was last modified.
Path	Path of the file.
Event Type	Type of file transfer event.
Event Time	The date and time when the file event occurred.

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 screenshot shows the 'File Transfer Logs' page with the 'Integration Logs' filter selected. The table displays 224 jobs. The columns are: Instance Name, Asset Type, Command, End Time, # of Files, and Status. The status column uses green checkmarks for 'Success' and red X marks for 'Failed'.

Instance Name	Asset Type	Command	End Time	# of Files	Status
Decrypt-558	File Transfer Task	File Transfer Task	Mar 6, 2020, 3:13 PM	0	Success
Decrypt-557	File Transfer Task	File Transfer Task	Mar 6, 2020, 3:10 PM	0	Failed
Decrypt-556	File Transfer Task	File Transfer Task	Mar 6, 2020, 3:09 PM	0	Failed
Encrypt-555	File Transfer Task	File Transfer Task	Mar 6, 2020, 3:06 PM	0	Success
Decrypt-554	File Transfer Task	File Transfer Task	Mar 6, 2020, 2:44 PM	0	Success
Encrypt-553	File Transfer Task	File Transfer Task	Mar 6, 2020, 2:43 PM	0	Success
Decrypt-552	File Transfer Task	File Transfer Task	Mar 6, 2020, 2:33 PM	0	Success
Decrypt-551	File Transfer Task	File Transfer Task	Mar 6, 2020, 2:32 PM	0	Failed
Decrypt-550	File Transfer Task	File Transfer Task	Mar 6, 2020, 2:32 PM	0	Failed
Encrypt-549	File Transfer Task	File Transfer Task	Mar 6, 2020, 2:29 PM	0	Success
PK1-548	API	Release	Mar 5, 2020, 2:35 PM	2	Success
spider_man-547	API	Release	Mar 4, 2020, 2:22 PM	2	Success
spider_man-546	API	Release	Mar 4, 2020, 2:18 PM	4	Success
spider_man-545	API	Release	Mar 4, 2020, 2:17 PM	4	Success
https_automation-544	API	Release	Mar 2, 2020, 1:52 PM	2	Success
https_automation-543	API	Collect	Mar 2, 2020, 1:46 PM	2	Success
https_automation-542	API	Collect	Mar 2, 2020, 12:34 PM	0	Failed

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, File Transfer Task, FTP Client, FTPs Client, or SFTP Client.
Command	Command type, such as Collect, Release, File Transfer Task, Upload, or Download.
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 Upload, Download, Release, Collect, or 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:



File Transfer Logs | Integration Logs

Jobs (224)

Instance Name	Asset Type	Command	End Time	# of Files	Status
Decrypt-558	File Transfer Task	File Transfer Task	Mar 6, 2020, 3:13 PM	0	Success
Decrypt-557	File Transfer Task	File Transfer Task	Mar 6, 2020, 3:10 PM	0	Success
Decrypt-556	File Transfer Task	File Transfer Task	Mar 6, 2020, 3:09 PM	0	Success
Encrypt-555	File Transfer Task	File Transfer Task	Mar 6, 2020, 3:06 PM	0	Success
Decrypt-554	File Transfer Task	File Transfer Task	Mar 6, 2020, 2:44 PM	0	Success
Encrypt-553	File Transfer Task	File Transfer Task	Mar 6, 2020, 2:43 PM	0	Success
Decrypt-552	File Transfer Task	File Transfer Task	Mar 6, 2020, 2:33 PM	0	Success
Decrypt-551	File Transfer Task	File Transfer Task	Mar 6, 2020, 2:32 PM	0	Failed
Decrypt-550	File Transfer Task	File Transfer Task	Mar 6, 2020, 2:32 PM	0	Failed
Encrypt-549	File Transfer Task	File Transfer Task	Mar 6, 2020, 2:29 PM	0	Success
PK1-548	API	Release	Mar 5, 2020, 2:35 PM	2	Success
spider_man-547	API	Release	Mar 4, 2020, 2:22 PM	2	Success
spider_man-546	API	Release	Mar 4, 2020, 2:18 PM	4	Success
spider_man-545	API	Release	Mar 4, 2020, 2:17 PM	4	Success
https_automation-544	API	Release	Mar 2, 2020, 1:52 PM	2	Success

1 - 25 of 224

< 1 of 9 >

Items per Page: 25

- To view the logs for Collect, Release, Upload, and Download API commands, click the instance name. To download the log of an instance, click **Download Job Log** in the Results area. Alternatively, click the **Actions** icon in the row that contains the instance and select **Download Job Log** on the **Integration Logs** page.

The following image shows the log of a sample Release command API:



PK1-548

Refresh

Runtime Environment: DXSuse12Sp2001

Started By: b2bp

Start Time: Mar 5, 2020, 2:35:44 PM

End Time: Mar 5, 2020, 2:35:55 PM

Duration: 00:00:12

Pattern: .*

File Server Username: PK1

Command: Release

Results

State: ✔ Success

Log: [Download Job Log](#)

Files Sent: 2

Error Message:

File Events (2)

Source File Path	Target File Path	File Size (KB)
------------------	------------------	----------------

Job Properties

The job properties display general properties about the instance.

The job properties include the following properties:

Property	Description
Job Name	The name of the job. This is for API, FTP Client, FTPs Client, and SFTP Client Asset Type.
Runtime Environment	Name of the runtime environment where you want to run the tasks.
Task Name	Name of the task. This applies only for the File Transfer Task Asset Type.
Task Type	Task type. This applies only for file ingestion task.
Connection	The connection used to send files.
Started By	The name of the user or schedule that started the job.
Start Time	Date and time when the job started.
End Time	Date and time when the job completed or stopped.
Duration	The amount of time the job ran before it completed or was stopped.
Pattern	The pattern of the transferred file.
Command	Command type

Results

The job results display the status of the job and error statistics.

The job results include the following properties:

Property	Description
State	Job status. A job can have one of the following statuses: <ul style="list-style-type: none"> - Running. The job is still running. - Success. The job completed successfully. - Failed. The job did not complete because it encountered errors - Aborted. The job was aborted.
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
Files sent	The number of files that are transferred, downloaded or uploaded to the target.
Error Message	Error message, if any, that is associated with the job.

File Events

This section shows the total number of files that are transferred, with information about each file.

The File Events section is updated each time a file is transferred, 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
Source File Path	The path from where the files are transferred.
Target File Path	The path to where the files are transferred.
File Size	The size of the file in bytes.
State	<p>The state of the file transfer. A file can have one of the following states:</p> <ul style="list-style-type: none">- Success. The file transfer completed successfully.- Failed. The file transfer did not complete because it encountered errors.- Processing*. The file transfer is still running.- Duplicate*. The task previously transferred a file with the same name, directory location, and size.- In Doubt*. The previous task instance encountered errors while transferring the file. Applicable for tasks where the source is configured to skip duplicate files. <p>* These states are available only for the Download command.</p> <p>You can monitor the State property to track the progress of the file transfer of each file.</p>

CHAPTER 7

Monitoring Data Ingestion and Replication jobs

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

Monitor

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

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

Operational Insights

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

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

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

Data Integration

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

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

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

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

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

Monitoring your ingestion and replication jobs

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

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

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

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

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

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

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

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

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

For descriptions of the columns in Monitor and Operational Insights, see [“Job properties” on page 101](#).

Monitoring all ingestion and replication jobs

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

Monitoring jobs in Operational Insights


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

The Data Ingestion and Replication page has the following tabs:

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

Overview tab

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




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

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

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

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




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

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

All Jobs tab

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



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

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


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

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

Monitoring jobs in Monitor

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

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



Instance Name	Asset Type	Location	Subtasks	Start Time	End Time	Rows Processed	Status
mssql_sf2_stop_job-73933	Database Ingestion and Re...	Default		Jun 6, 2024, 4:48 AM	Jun 6, 2024, 4:54 AM	5	Undeployed
mssql_sf2_stop_job-73932	Database Ingestion and Re...	Default		Jun 6, 2024, 4:48 AM	Jun 6, 2024, 4:54 AM	5	Undeployed
oracle_spispe_unload_060624-73934	Database Ingestion and Re...	Pruthvi/Oracle		Jun 6, 2024, 4:44 AM	Jun 6, 2024, 4:46 AM	1933	Completed
ppsql_unload_6-73935	Database Ingestion and Re...	dimeg/2023.07.M...		Jun 6, 2024, 4:43 AM	Jun 6, 2024, 4:44 AM	13	Completed
ppsql_test_perm1-73931	Database Ingestion and Re...	dimeg/2023.07.M...		Jun 6, 2024, 4:39 AM	Jun 6, 2024, 4:40 AM	13	Completed

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

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

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

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

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

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

Job properties

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

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

Property	Description
Instance Name	<p>The generated name of the job instance in the following format: <code><taskname>-<job_instance_number></code></p> <p>You can click the instance name to view detailed information about the job. Note: If you edit the name of the associated ingestion and replication task, the job name remains the same.</p>
Location	<p>The project or project\subfolder, where the task definition associated with the job exists. For example: <code>Myproject\Oracle</code></p> <p>If you move a task definition to another folder, the Location value is not updated.</p>
Task Type	<p>The type of ingestion and replication task. The type can be one of the following values:</p> <ul style="list-style-type: none"> - Application Ingestion and Replication Task - Database Ingestion and Replication Task - File Ingestion and Replication Task - Streaming Ingestion and Replication Task
Runtime Environment	<p>The name of the runtime environment that includes the Secure Agent on which the job runs.</p>
Start Time	<p>For application ingestion and replication and database ingestion and replication jobs, the date and time when the job was deployed.</p> <p>For file ingestion and replication jobs, the date and time when the job started.</p> <p>For streaming ingestion and replication jobs, the date and time when the job was deployed.</p>
Duration	<p>For application ingestion and replication and database ingestion and replication jobs, the amount of time that the job has run since it was deployed. For jobs that have the Completed, Stopped, Failed, or Aborted status, the amount of time between the date and time the job was deployed and when it acquired its current status.</p> <p>For file ingestion and replication jobs, the amount of time that the job has run.</p> <p>For streaming ingestion and replication jobs, the amount of time that the job has been running.</p>
Status	<p>The current status of the job, such as Deploying, Up and Running, or Undeployed.</p> <p>The set of valid statuses vary by type of ingestion and replication task. For more information, see the "Job Overview tab" section in "Application ingestion and replication job details" on page 103, "Database ingestion and replication job details" on page 109, "File ingestion and replication job details" on page 116 ("Results" section), or "Streaming ingestion and replication job details" on page 119.</p>

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

Property	Description
Instance Name	<p>The generated name of the job instance in the following format: <taskname>-<job_instance_number></p> <p>You can click the instance name to view detailed information about the job.</p> <p>Note:</p> <ul style="list-style-type: none"> - If you edit the name of the associated ingestion and replication task, the job name remains the same. - The job name contains an underscore (_) instead of a hyphen (-): <taskname>_<job_instance_number>
Location	<p>The project or project\subfolder, where the task definition associated with the job exists. For example: Myproject\Oracle</p> <p>If you move a task definition to another folder, the Location value is not updated.</p>
Subtasks	<p>The number of subtasks on source objects.</p> <p>Note: For application ingestion and replication jobs and database ingestion and replication jobs, the subtask count is not displayed.</p>
Start Time	<p>For application ingestion and replication and database ingestion and replication jobs, the date and time when the job was deployed.</p> <p>For file ingestion and replication jobs, the date and time when the job started.</p> <p>For streaming ingestion and replication jobs, the date and time when the job was deployed.</p>
End Time	<p>For application ingestion and replication and database ingestion and replication jobs, the date and time when the job ended because it completed processing, was stopped, or failed. This field is not displayed for running jobs.</p>
Rows Processed	<p>For application ingestion and replication jobs and database ingestion and replication jobs, the total number of rows that the job has processed and successfully written to the target. If a job encounters an error during row processing, the job status switches to Failed and the rows with the error are not included in the Rows Processed count.</p> <p>For file ingestion and replication jobs, the total number of files that the job has processed and successfully written to the target. If a job encounters an error during file processing, the job status switches to Failed and the files with the error are added in the Failure Rows count. In this case, the Rows Processed property displays the total number of files processed irrespective of their statuses.</p> <p>For streaming ingestion and replication jobs, the total number of events that the job has processed and successfully written to the target. If a job encounters an error during event processing, the job status switches to Running with Warning or Running with Error and the events associated with the error are not included in the Rows Processed count.</p>
Status	<p>The current status of the job, such as Deploying, Up and Running, or Undeployed.</p> <p>The set of valid statuses vary by type of ingestion and replication task. For more information, see the "Job Overview tab" section in "Application ingestion and replication job details" on page 103, "Database ingestion and replication job details" on page 109, "File ingestion and replication job details" on page 116 ("Results" section), or "Streaming ingestion and replication job details" on page 119.</p>

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

Viewing details for an ingestion and replication job

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

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

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

Application ingestion and replication job details

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

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

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

Task Summary

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

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

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

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

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

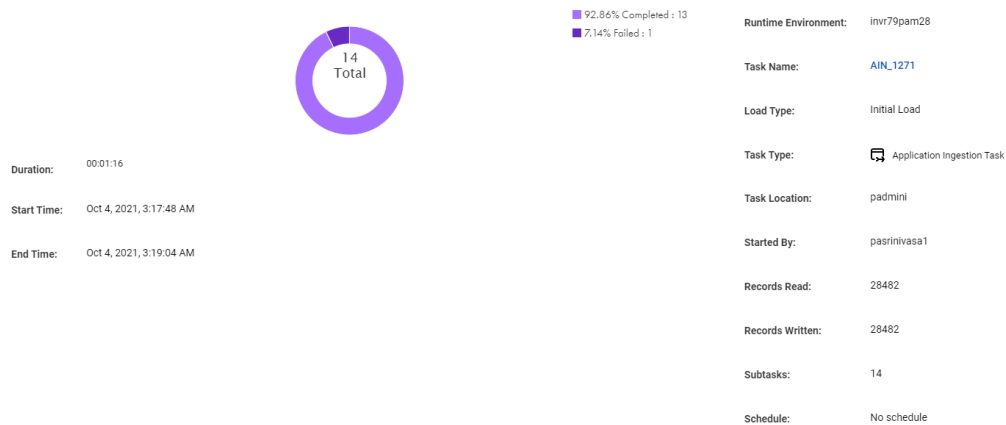
Note: For initial load jobs, you can get the job log for a specific source object from the **Object Detail** tab.

Overview

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

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

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



The following table describes the summary job properties and statistics:

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

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

Object Detail

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

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

Object Detail (14)

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

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

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

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

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

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

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

Alerts

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

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

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

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

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

The list of alerts includes the following columns:

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

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

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

Database ingestion and replication job details

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

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

The following image displays the panes:

The screenshot displays the job details for 'ora_gen2_combined_73735'. The 'Task Summary' pane shows the job is 'Running with Warning'. The 'Overview' pane shows a donut chart with '12 Total' objects. The 'Object Detail' pane shows a table with the following data:

Object	Target Object	Agent Name	Stage	Status	Log
SREE_SRC_ALLCHAR_BK	SREE.com_SRC_ALLCHAR_BK	Informatica.com	Normal	Running	Select log
SREE_SRC_ALDTE_BK	SREE.com_SRC_ALDTE_BK	Informatica.com	Normal	Running	Select log

The 'Alerts' pane shows 13 alerts, including one Critical alert (DBMIP_27008) and several Info alerts (DBMIP_27002).

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

Task Summary

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

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

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

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

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

Overview

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

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

The following table describes the summary job properties and statistics:

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

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

Object Detail

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

The following image shows a sample expanded object row:

Object	Target Object	Status	Log
▼ DBMICDC_pgs013_rc_rongetype1_mri	KARTHEEK.pgs013_rc_rongetype1_mri_uno=flake	Stopped	Select log
Inserts	Updates	Deletes	LOBs
11	4	11	0
▶ DBMICDC_pgs014_rc_rongetype2_mri	KARTHEEK.pgs014_rc_rongetype2_mri_uno=flake	Stopped	Select log

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

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

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

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

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

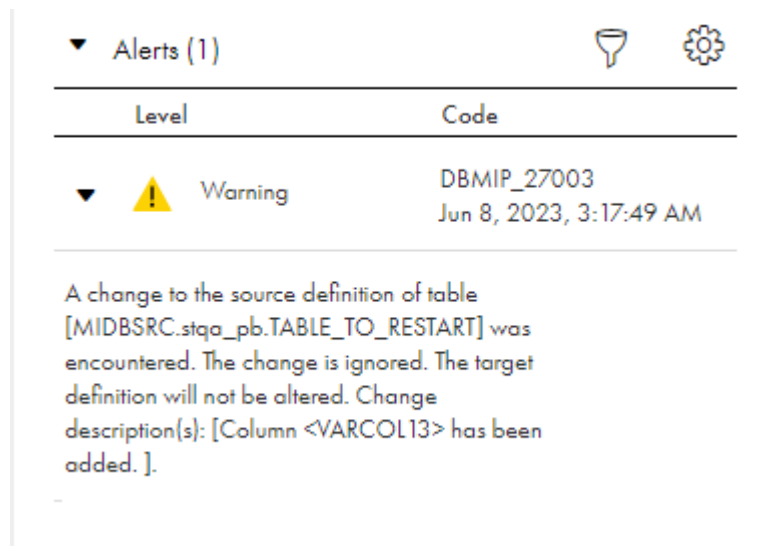
Column	Description
CDV Last Run	<p>If data validation is available for the task, this column shows the date and time when the data validation was last performed for the given subtask. If the data validation has not been run, the field value is N/A.</p> <p>This column appears in initial load jobs that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

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

Alerts

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

The following image shows this pane with an alert example:



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

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

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

The list of alerts includes the following columns:

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

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

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

File ingestion and replication job details

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

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

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

The screenshot displays the details of a file ingestion and replication job. The job name is `adlsgen2_to_ftps_1SatJun122021094156-124505`. The job properties show it was started by `filelistener` on `Jun 12, 2021, 10:31:16 AM` and ended on `Jun 12, 2021, 10:31:31 AM` with a duration of `00:00:15`. The results section indicates the job was successful, with 14 success files, 0 error files, and 0 duplicate files. The error message is "Job completed normally".

The File Events table shows the following data:

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

Job Properties

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

The following table describes the job properties:

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

Results

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

The following table describes the job results:

Property	Description
State	Job status. A job can have one of the following statuses: <ul style="list-style-type: none">- Running. The job is still running.- Success. The job completed successfully.- Failed. The job did not complete because it encountered errors- Aborted. The job was aborted. <p>Note: When the Secure Agent is unavailable while running a file ingestion job, the job remains in an unresponsive state, and after 200 minutes, its status changes to Aborted.</p>
Session Log	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/MassIngestionRuntime/data/taskLogs
Success Files	Number of files that are successfully transferred, downloaded, and uploaded to the target.
Error Files	Number of files that were not transferred to the target.
Duplicate Files	Number of files that were identified as duplicates.
Error Message	Error message, if any, that is associated with the job.

File Events

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

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

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

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

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

Streaming ingestion and replication job details

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

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

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

Overview tab

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

The following table describes the job overview properties:

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

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

Alert tab

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

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




The following table describes the job alert properties:

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

Performance tab

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

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



The following table describes the job performance properties:

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

Past Run tab

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

The following table describes the properties displayed for each run:

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

Data Ingestion and Replication alerts

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

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

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

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

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

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

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

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

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

Configuring alerts for Data Ingestion and Replication jobs

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

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


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

CHAPTER 8

Monitoring advanced clusters

View information about advanced clusters on the **Advanced Clusters** page. Advanced clusters that use CLAIRE report on the estimated cloud infrastructure savings due to CLAIRE's intelligent optimizations, key areas that contribute to those savings, and additional financial insights to help you govern and monitor cloud spending and enable FinOps capabilities.

The following image shows the **Advanced Clusters** page:



The **Advanced Clusters** page includes the following panels:

CLAIRE insights

If CLAIRE recommendations are enabled in your organization, you can view CLAIRE insights for advanced clusters that use a CLAIRE-powered configuration. CLAIRE insights show the estimated cloud infrastructure savings due to CLAIRE's intelligent optimizations.

You can view the average savings per hour, the total savings, and a graph that visualizes the key areas that contribute to infrastructure savings. The graph divides savings into the following key areas:

- Spot Instance selection. Savings gained by using Spot Instances over On-Demand Instances where appropriate. CLAIRE selects Spot Instances only in clusters that are optimized for cost.
- Instance sizing. Savings gained by selecting instance types based on the cluster optimization preference and the resources that the cluster uses to run a typical workload.
- Cluster scaling. Savings gained by scaling the cluster and local storage based on the expected workload.

- Job scheduling. Savings gained by scheduling jobs to efficiently use cluster resources.
- Smart shutdown. Savings gained by shutting down the cluster when the cluster is expected to be idle

You can also view mapping statistics to understand how many mappings and jobs have been created in your organization.

Advanced Clusters table

Lists all advanced clusters in the organization. You can click each status above the table to quickly filter the table by cluster status, and you can drill down on a cluster to view more detailed information.

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

```
<advanced configuration>-<runtime environment>
```

For information about cluster statuses, see [“Cluster statuses” on page 126](#).

To view detailed information about an individual cluster, you can drill down on the cluster and its cluster instances.

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, you can click the "Updates Available" message or the **Refresh** icon. You can also enable **Auto Refresh** to automatically refresh the information every 20 seconds.

Information can become out-of-date when the state of a cluster changes, such as when you run a job to start a cluster or worker nodes are added to the cluster.


Note: Auto Refresh only refreshes the cluster states. To refresh CLAIRE insights, you must refresh the page manually.

To monitor clusters on the **Advanced 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 advanced configurations to the role.

Monitoring an individual advanced cluster

Monitor an individual advanced cluster to view CLAIRE insights for the cluster, the infrastructure costs to run the cluster, and a list of cluster instances.

The following image shows the details that you can view for an individual cluster:



You can view the following details:

CLAIRE insights

If the advanced cluster uses a CLAIRE-powered configuration, you can view CLAIRE insights for the cluster. You can also view a summary of the CLAIRE recommendations that are available. To view the recommendations, navigate to the advanced configuration for the cluster in Administrator.

Infrastructure costs over time

If the advanced cluster uses a CLAIRE-powered configuration, you can view a graph of the infrastructure costs over time. For each time interval, you can view the actual infrastructure costs and compare them to your budget.

Cluster instances

You can view a list of cluster instances that are associated with the advanced cluster. Drill down on a cluster instance to view its activity log, lifecycle graph, configuration, as well as the jobs that were submitted to the cluster instance.

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

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, you can click the "Updates Available" message or the **Refresh** icon. You can also enable **Auto Refresh** to automatically refresh the information every 20 seconds.

Information can become out-of-date when the state of a cluster changes, such as when you run a job to start a cluster or worker nodes are added to the cluster.

Cluster statuses

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

The following table describes cluster statuses:

Status	Description
Starting	The cluster is starting. A cluster starts as soon as you run a job.
Running	The cluster is running and processing jobs.
Stopping	The cluster is stopping. The jobs that were running on the cluster have completed and the cluster has reached the idle timeout in the advanced 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 a 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 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 advanced configuration, choosing to disable the advanced 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 an advanced cluster instance

Monitor an advanced cluster instance to view its activity log, lifecycle graph, advanced configuration, and the jobs that were submitted to the cluster instance.

You can drill down on each cluster instance to view 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 advanced 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 advanced configuration that a Secure Agent uses to create the advanced cluster. To edit the advanced configuration, you must use Administrator.

Jobs

A list of all jobs that were submitted to the cluster. You can stop and restart jobs and download log files.

When you monitor 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.

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 **Advanced 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 advanced 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 a 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 advanced 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. If the Secure Agent stops unexpectedly and restarts on a different machine, the agent can recover the cluster only if the version of the Elastic Server on the new agent is the same as the version of the Elastic Server on the previous agent.

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 **Advanced 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:



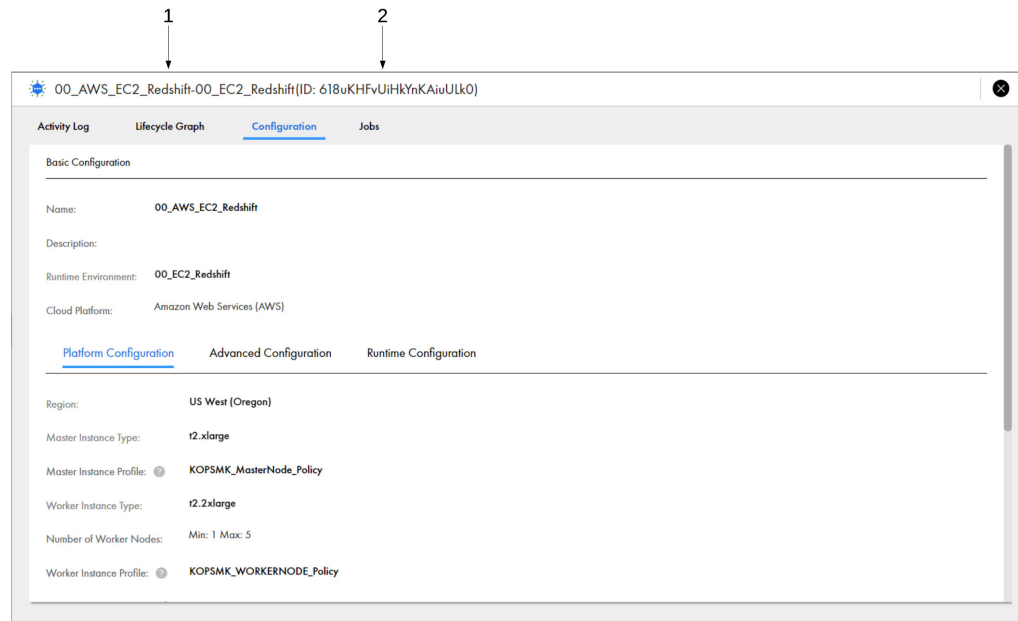
1. Cluster name
2. Cluster ID
3. Time range

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 **Advanced Clusters** page. The configuration that you view is the advanced configuration that you use to provision resources for an advanced 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 jobs that were submitted to a cluster on the **Jobs** tab after you drill down on the cluster instance from the **Advanced Clusters** page. The **Jobs** tab lists the jobs that are currently running and the jobs that have completed.

Use the **Jobs** tab to analyze job failures and debug both jobs and the advanced cluster.

To avoid unnecessary failures, check the status of the advanced cluster before you run a 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 advanced cluster subtasks" on page 28](#).

CHAPTER 9

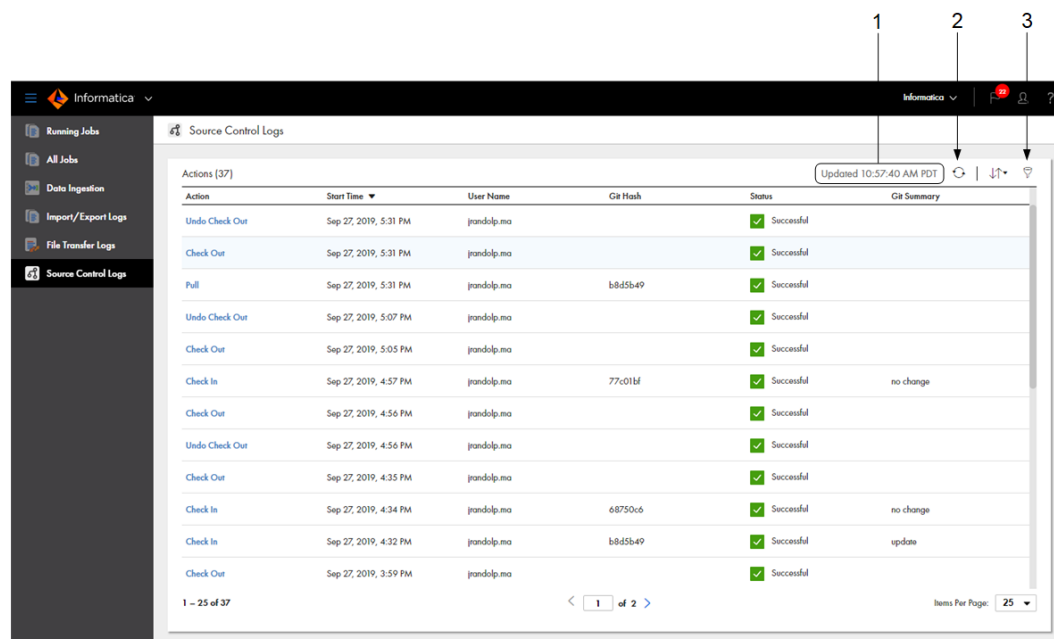
Monitoring source control logs

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 30 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.

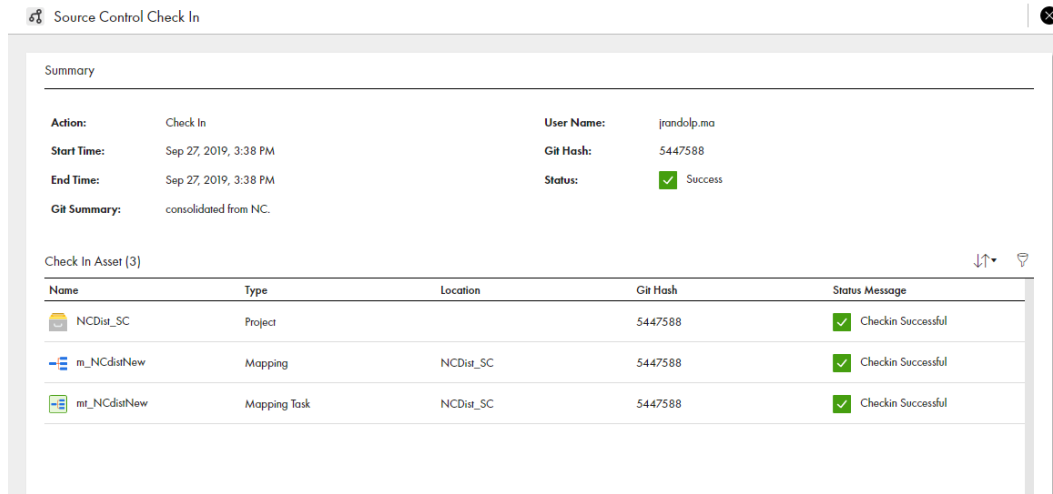
The action properties include the following properties:

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: <ul style="list-style-type: none">- 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: <ul style="list-style-type: none">- Check in- Pull
Status	Status of the action. An action can have one of the following statuses: <ul style="list-style-type: none">- 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

A

- advanced cluster
 - viewing job results [29](#)
 - viewing reprocessing details [26](#)
- advanced clusters
 - cluster instance [126](#)
 - cluster status [126](#)
 - monitoring [123](#)
 - monitoring cluster jobs [130](#)
 - monitoring individual clusters [125](#)
 - monitoring the cluster activity log [127](#)
 - viewing the advanced configuration [130](#)
 - viewing the lifecycle graph [129](#)
- advanced mode
 - CLAIRE recommendations [27](#)
- alerts
 - application ingestion and replication jobs [121](#)
 - database ingestion and replication jobs [121](#)
- AS2 file transfer
 - AS2 log properties [79](#)
 - job details [80](#)
- AS2 file transfers [79](#)

C

- Cloud Application Integration community
 - URL [7](#)
- Cloud Developer community
 - URL [7](#)
- code tasks
 - monitoring [31](#)
 - viewing Spark application task details [34](#)

D

- Data Ingestion and Replication jobs
 - alerts
 - configuring for application or database ingestion and replication jobs [122](#)
 - configuring alerts [122](#)
- Data Integration community
 - URL [7](#)
- Data Profiling jobs
 - monitoring [70](#)

E

- elastic clusters
 - cluster events [128](#)
- elastic mappings
 - viewing incremental file load details [26](#)

- error rows files
 - synchronization tasks [65](#)
 - viewing [67](#)
- export files
 - downloading [77](#)
- exports
 - customizing properties [75](#)
 - downloading an export file [74, 76](#)
 - downloading export files [77](#)
 - downloading log files [78](#)
 - filtering [75](#)
 - monitoring [74](#)
 - properties [74](#)
 - sorting [75](#)
 - viewing details [76](#)

F

- file transfer jobs
 - viewing details [80, 85, 89](#)
- file transfer log
 - properties [79, 84, 86](#)
- file transfer logs [79, 83, 86, 87](#)

H

- HTTPS file transfer
 - HTTPS log properties [86](#)
- HTTPS file transfers [86](#)

I

- imports
 - customizing properties [75](#)
 - downloading log files [78](#)
 - filtering [75](#)
 - monitoring [74](#)
 - properties [74](#)
 - sorting [75](#)
 - viewing details [76](#)
- InfraSpark0
 - viewing Spark task details [30](#)
- Informatica Global Customer Support
 - contact information [8](#)
- Informatica Intelligent Cloud Services
 - web site [7](#)

J

- jobs
 - AS2 log properties [79](#)
 - customizing properties [16](#)

jobs (continued)

- downloading log files [67, 73](#)
- error rows files [65](#)
- filtering [16](#)
- HTTPS log properties [86](#)
- log files [65](#)
- logs for advanced clusters [68](#)
- monitoring [9](#)
- monitoring all jobs [9](#)
- monitoring code tasks [31](#)
- monitoring masking tasks [40](#)
- monitoring running jobs [10](#)
- monitoring your jobs [11](#)
- naming convention [70](#)
- properties [14, 70](#)
- restarting [59, 72](#)
- SFTP log properties [84](#)
- sorting [16](#)
- stopping [59, 72](#)
- viewing detailed information [17](#)
- viewing details for linear taskflows [54](#)
- viewing details for mappings [20](#)
- viewing details for taskflows [45](#)
- viewing details for tasks [20](#)
- viewing error rows files [67](#)
- viewing incremental file load details [26](#)
- viewing individual task results for linear taskflows [55](#)
- viewing job properties [20](#)
- viewing job properties for linear taskflows [54](#)
- viewing job results [21, 29](#)
- viewing job results for linear taskflows [55](#)
- viewing mapping task details [22](#)
- viewing properties [14](#)
- viewing replication task details [43](#)
- viewing reprocessing details [26](#)
- viewing Spark application task details [34](#)
- viewing Spark task details [30](#)
- viewing subtasks [12](#)

L

linear taskflows

- viewing individual task results [55](#)
- viewing job details [54](#)
- viewing job properties [54](#)
- viewing job results [55](#)

log files

- downloading [67, 73](#)
- failed jobs [65](#)

M

maintenance outages [8](#)

mappings

- viewing job details [20](#)
- viewing job properties [20](#)
- viewing job results [21](#)
- viewing mapping task details [22](#)
- viewing replication task details [43](#)
- viewing Spark task details [30](#)

MLLP file transfer

- event details [89](#)

MLLP file transfers [87](#)

monitoring

- all jobs [9](#)
- AS2 file transfers [79](#)

monitoring (continued)

- Data Profiling jobs [70](#)
- HTTPS file transfers [86](#)
- imports and exports [74](#)
- jobs [9](#)
- MLLP file transfers [87](#)
- running jobs [10](#)
- SFTP file transfers [83](#)
- your jobs [11](#)
- monitoring ingestion and replication jobs
 - database ingestion and replication job details [109](#)
 - Job Overview tab [109](#)
 - job properties in job lists [101](#)
 - monitoring all jobs from Operational Insights and Monitor [98](#)
 - monitoring My Jobs [97](#)
 - Object Details tab [109](#)
 - overview [96](#)
 - viewing job details [103](#)
- monitoring source control logs [132](#)
- My Jobs
 - view subtask details [71](#)

S

SFTP file transfer

- event details [85](#)
- SFTP log properties [84](#)

SFTP file transfers [83](#)

source control action details [134](#)

status

- Informatica Intelligent Cloud Services [8](#)

streaming ingestion and replication jobs [119](#)

system status [8](#)

T

taskflows

- monitoring [9](#)
- monitoring all taskflows [9](#)
- monitoring running taskflows [10](#)
- monitoring your taskflows [11](#)
- resuming [60](#)
- suspending [60](#)
- viewing job details [45](#)

tasks

- monitoring [9](#)
- monitoring all tasks [9](#)
- monitoring running tasks [10](#)
- monitoring your tasks [11](#)
- viewing incremental file load details [26](#)
- viewing job details [20](#)
- viewing job properties [20](#)
- viewing job results [21](#)
- viewing mapping task details [22](#)
- viewing replication task details [43](#)
- viewing reprocessing details [26](#)
- viewing Spark task details [30](#)

trust site

- description [8](#)

U

upgrade notifications [8](#)

V

viewing source control action details [134](#)

W

web site [7](#)