SAS® Job Monitor 2.2

2.2

This document might apply to additional versions of the software. Open this document in SAS Help Center and click on the version in the banner to see all available versions.

What's New for SAS Job Monitor 2.2 . . . 1

About SAS Job Monitor . . . 2
- Overview . . . 2
- SAS Job Monitor for SAS Data Integration Studio . . . 2
- SAS Job Monitor for DataFlux Data Management Server . . . 2

Configuring SAS Job Monitor . . . 3
- SAS Data Integration Studio . . . 3
- DataFlux Data Management Server . . . 4
- Additional Servers . . . 5

Using the Monitoring Center . . . 6
- Overview of the Monitoring Center . . . 6
- Job Drill-In . . . 8

What's New for SAS Job Monitor 2.2

SAS Job Monitor components have been modernized for performance and security. The features are otherwise unchanged.

About SAS Job Monitor

Overview

SAS Job Monitor is a component of SAS Environment Manager that integrates information from SAS Data Integration Studio, DataFlux Data Management Server, and specific jobs from DataFlux Data Management Studio. It enables you to oversee the state of jobs that are run from these applications.

SAS Job Monitor reads job logs at specified locations and displays run-time statistics from the logs, enabling you to monitor status information and performance statistics. Using the jobs table, you can view historical run times and also drill into a job for details, trends, and status. Since the information in the jobs table is parsed from job logs of SAS Data Integration Studio and DataFlux Data Management Server, minor configuration changes are required for these platforms. The information for DataFlux Data Management Server also applies to jobs run in DataFlux Data Management Studio using the **DMPEXEC** tool.

Note: Migration is not supported for SAS Job Monitor. Old monitoring information is considered transient and is not stored. The user is expected to redeploy jobs and scheduled flows after migrating to a new release version.

SAS Job Monitor for SAS Data Integration Studio

Job monitoring information for SAS Data Integration Studio is taken from directories where SAS Data Integration Studio batch job logs reside. The logs are parsed for key information about the jobs as they are written. See “SAS Data Integration Studio” on page 3 for additional information.

SAS Job Monitor for DataFlux Data Management Server

Job monitoring for DataFlux Data Management Server is based on the same principles as monitoring for SAS Data Integration Studio. Directories where the server logs reside are monitored for activity. Information taken from the logs is written and displayed in the Monitoring Center. Because DataFlux Data Management Studio does not generate a unique log for each job, you cannot use SAS Job Monitor to display run-time statistics for jobs that are executed interactively in DataFlux Data Management Studio. The jobs must be executed with the **DMPEXEC** command or passed to a DataFlux Data Management Server.

Job runs are monitored by extracting information from the log files that are written as jobs are running. Therefore, jobs must be producing logs and SAS Job Monitor must be pointed to one or more directories where these logs reside. Logs in subdirectories are discovered automatically. By

default, job logs do not contain run-time nodes' statistics. In that case, SAS Job Monitor retrieves job-level details and the final statistics for all nodes at the end of the job run. See “DataFlux Data Management Server” on page 4 for additional information.

Configuring SAS Job Monitor

SAS Data Integration Studio

Configuration Requirements

Job statistics collection for SAS Data Integration Studio is dependent on automatically generated comments for jobs and steps, and information captured from ARM messages that are generated in the job log. The job must have **Collect Runtime Statistics** enabled to generate ARM messages. In addition, **Collect Table Statistics** must be active so that the **Rows Processed** field in the monitoring center is populated with data. For more information about the setup that is required to monitor jobs that were created in SAS Data Integration Studio, see the topic “Meeting Prerequisites for Collecting Job Statistics” in the *SAS Data Integration Studio: User’s Guide*. This topic appears in the “Managing Jobs” chapter.

Environment Variables

Locale

The locale in which SAS Data Integration Studio is running determines the strings that are used to generate job and step header comments. SAS Job Monitor contains a set of locale-specific strings that it uses to find these strings in the headers so that values can be extracted. Examples of these strings are **Step Name** or **Job Path**. The locale in which the agent for SAS Job Monitor is running determines which values are used to parse the header comments, so it must match the locale for SAS Data Integration Studio. The locale that the agent uses can be overridden by changing the values on the SAS Data Integration configuration page. In addition, the agent for Job Monitor expects the log files to be written in UTF-8. If the log file is written in an encoding other than UTF-8, you must change the encoding on the SAS Data Integration Server configuration page to match.

To change the locale, enter an encoding value in the **jobmonitoring.encoding** field. Here is an example:

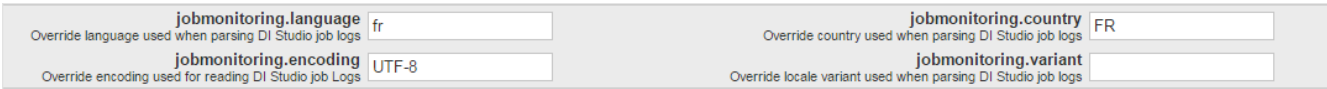

Sometimes it is necessary to also change the language and country settings to match the environment in SAS Data Integration Studio. Using the Data Integration Server configuration page, change these environment variables using **jobmonitoring.language**, **jobmonitoring.country**, and **jobmonitoring.variant** fields.

DataFlux Data Management Server

Configuration Requirements

Reporting for SAS Job Monitor is dependent on the **DF.Monitor** logger, which is enabled by default. All messages are written at the INFO level. To change the frequency of statistics collection for DataFlux Data Management Server or DataFlux Data Management Studio, the BASE/MONITOR_FREQUENCY parameter must be added and configured in `app.cfg`. For more information about the setup that is required to monitor jobs that were created in DataFlux Data Management Studio, see the topic “Displaying Run-Time Statistics in SAS Job Monitor” in the *DataFlux Data Management Studio: User’s Guide*. This topic appears in the “Maintaining Data Jobs” chapter and the “Maintaining Process Jobs” chapter.

Note: If BASE/MONITOR_FREQUENCY is set too low, then there can be a slight inconsistency between the number of loops a job executes during its run and the number of executions actually captured in the job log and displayed in the monitoring center. When monitor frequency is set at a higher rate, more snapshots are taken, which shows more loops or executions captured in the job log. However, it should be noted that using a higher monitor frequency could impact job performance.

Environment Variables

Locale

The agent for SAS Job Monitor expects the log files to be written in UTF-8. If the log file is written in a different encoding, you must change the encoding used in SAS Job Monitor to match using the server configuration page. Another option is to change the encoding that is used to write the logs by updating the configuration on DataFlux Data Management Server using the BASE/JOB_LOG_ENCODING option in the server’s configuration file.

This also applies to jobs that are run using SAS Visual Process Orchestration. For example, suppose that you execute an Orchestration job on a server, and the server’s locale setting results in a job log that is not in UTF-8. You might need to update some default options for SAS Job Monitor, or it might not be able to read the log. This situation is most likely to occur for locales that do not use the Western European encoding. To change the encoding, use the **jobmonitoring.encoding** override option in the configuration:

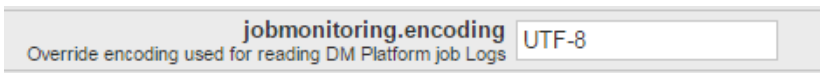

Additional Servers

Adding Servers for Job Monitoring

Adding a Server from Auto-Discovery

During installation, servers are discovered and populated in the Auto-Discovery portlet under the **Dashboard** tab. The server is prefixed with a host or platform name.

1. Using the Auto-Discovery portlet, select an item under **Resource Name** and click **Add to Inventory**. The message “No resources to display” appears if there are no items for auto-discovery.
2. **Configuration Properties** is displayed so that you can configure the new server. See “Configuring a Server” on page 5 for more details.

Adding a Server Manually

Most servers are auto-discovered but you can add a server manually using the **Resources** tab.

To add a server for monitoring:

1. Select the **Resources** tab and click **Browse** to produce a list of platforms.
2. Drill into the desired platform and select **Tools Menu**, **New Server** from the drop-down list.
3. At the **New Server** page, under **General Properties**, enter a *Name* and a *Description*.
4. Under **Type and Host Properties** select a server type from the drop-down list. If you are adding a Visual Process Orchestration server, select **SAS DM Platform 2.5+** as the server type.
5. Enter the full installation path for the server and click **OK**.

Immediately after adding a server, you are prompted to configure the server. Select this link to configure the new server. You can also scroll to the bottom of the page and select **Edit** at the **Configuration Properties** section. See “Configure Batch Jobs for Monitoring” on page 5 for configuration procedures.

Configuring a Server

Configure Batch Jobs for Monitoring

SAS Job Monitor is disabled by default. The **Configuration Properties** section contains the necessary options to enable monitoring.

Normally, **Configuration Properties** is presented at the time you add a new server. To navigate to **Configuration Properties** at any other time, do the following:

1. On the platform’s resources page, select **Inventory** and open the new server item.
2. Select **Tools Menu** and **Configure Server** from the drop-down list. **Configuration Properties** is displayed.

Figure 1 *Job Monitor Configuration Properties*

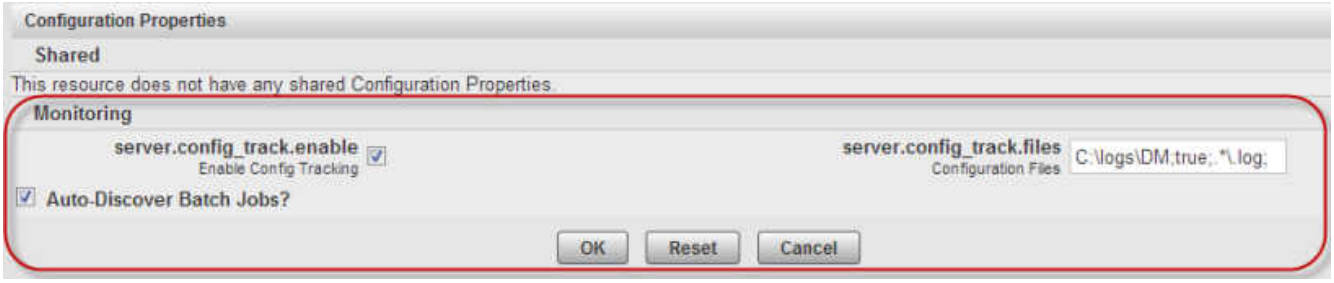

To enable monitoring, set the following configurations:

- For **server.config_track.enable**, select the **Enable Config Tracking** check box to enable monitoring of job logs.
- For **server.config_track.files**, : in the **Configuration Files** field, specify the directory or directories to be monitored. If you specify multiple directories, separate each entry with a comma. Each directory entry consists of the following three parts separated by semicolons:
  - the full path to the directory to be monitored
  - *true* or *false* to indicate whether subdirectories should be monitored recursively,
  - a file name filter that is a regular expression. Only files that match the filter are treated as log files.

  One example, `C:\DMServer\logs;true;.*/.log;`
- Select the **Auto-Discover Batch Jobs?** check box.
- Click **OK** to save the configuration.

Using the Monitoring Center

Overview of the Monitoring Center

Overview

The Monitoring Center for SAS Job Monitor is accessed through the toolbar in SAS Environment Manager. After launching the Monitoring Center, a default job table is displayed, which reflects all of the job types from all of the platforms and includes all states within the last eight days. You can sort jobs by applying filters.

Launching the Monitoring Center

Access the Monitoring Center using the toolbar in SAS Environment Manager:

1. Select **Analyze**.

2 Select **Monitoring Center** from the drop-down menu.

Applying Filters

Using the Date Range Selector

The jobs in the jobs table are filtered according to the date range reflected in the date range selector. You can change the date range by using the drop-down list or selecting **Edit Range** to specify a time frame:

Using the Job Filter

Use the job filter to limit the jobs that are displayed. You can specify a job name or, using drop-down menus, you can select a job type, platform, or state. You can also filter jobs by User ID.

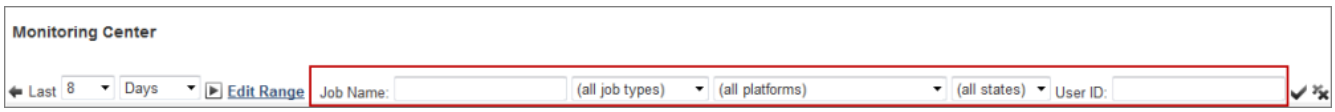

1 Use one of the following categories to filter jobs:

- Specify a job name in the **Job Name** field to filter by that name.
- Select a job type to display a list of jobs of that type. The default is **(all job types)**.
- Select a platform to display the list of jobs restricted to the selected platform. The default is **(all platforms)**.
- Select a state to display jobs with a particular state. The choices are **Running**, **Completed**, **Error**, **Stopped**, or **Warning**. The default is **(all states)**.
- Specify a user ID in the **User ID** field to display jobs run by a specific user.

2 Click ✔ to apply the filter.

About the Jobs Table

The job table displays the list of all the jobs that have run or are currently running. By default, the table is sorted by start time with the latest start time at the top. The following columns appear in the Jobs Table:

Job Name
includes an icon of the job type and the name of the job. A job can appear multiple times if it has run multiple times. Clicking on a job name drills into the details of the job.

Job Type
displays the type of job. The job types that are supported are as follows:

- SAS Data Integration Studio job
- DataFlux Data Management Server data job
- DataFlux Data Management Server process job

Platform
displays the name of the platform as defined by the system.

Status
displays the state of the job—Running, Error, Warning, or Complete, and is based on information parsed from the job log. A job has a status of running if it does not have an end time. A job does not appear as running until the first step has finished. If the job has finished, it reflects a status of Error, Warning, or Complete.

% Complete
displays the completed percentage compared to the mean of the previous number of runs. If the job has not completed, the maximum value shown is 99%, even if the job has taken longer than the N-run mean. If the job completes without errors, the value is set to 100%. If the job completes with errors and the N-run mean can be calculated, then the value is set to the percentage of the run time compared to the mean with a maximum value of 100%.

If there is no N-run mean, the cell remains blank until the job has finished running, at which time the value shows as 100.

Start Time
displays the start time of the job.

End Time
displays the time at which the job stopped.

Run Time
displays the total amount of time that the job took to run.

Run Time Alert
displays an alert icon when the difference of **Run Time** and **N-Run Mean** is greater than twice the standard deviation and if there are at least two previous executions.

20 Run Mean (N-Run Mean)
displays the mean of the previous 20 runs. This column is blank if there are no previous runs.

Trend Column
displays icons indicating the run-time trend.

User ID
displays the identification of the user running the job.

Job Drill-In

Overview

The jobs table displays the list of jobs that have run or are currently running. By default the table is sorted by start time with the latest start time displayed first. Clicking on a job in the Jobs table drills into the details of a SAS Data Integration Studio or a DataFlux Data Management Server job.

Note: When drilling into a job or its steps, use the **Jobs >** link to return to the jobs table. Do not use your browser's back button.

Each job is categorized as shown here:

- Steps
- Job History
- Job Trend
- Status Messages
- Log

Steps

The steps tab shows a table listing the steps within the job. By default the table is sorted by start time with the latest start time at the top. Select the job's step in the steps table to drill into the details of the step. A new screen is displayed that shows four tabs: **Step History**, **Step Trend**, **Status Messages**, and **Job Log**.

Note: If the step is a container for child steps, the Steps drill-in reflects a fifth **Steps** tab, which lists the inside steps that can also be drilled into.

- The **Step History** tab shows a table that lists the historical runs of the step. By default the table should be sorted by start time with the latest start time at the top. The step run that has been drilled into is highlighted in the list.
- The **Step Trend** tab shows a column chart that reflects the steps run time over time. The step run that has been drilled into is highlighted in the graph. The date range of the graph can be changed using the drop-down menus, or by selecting **Edit Range** to set a date range with specific dates and times.
- The **Status Messages** tab shows the status for each step
- The **Log** tab displays the log for the job containing the selected step.

Job History

The **Job History** tab includes a table that lists the historical runs of the job. By default the table should be sorted by start time with the latest start time at the top. The job run that has been drilled into is highlighted in the list. Although it does not include all the columns in the jobs list, the columns mostly behave the same as the columns in the jobs list. You can change the date range using the drop-down menus or by selecting **Edit Range** to set a date range with specific dates and times.

Job Trend

The **Job Trend** tab shows a stacked column chart that shows the job's run time over time. The job run that has been drilled into is highlighted in the graph. The columns consist of a stack of the steps that ran in the job. You can change the date range using the drop-down menus or by selecting **Edit Range** to set a date range with specific dates and times.

Status Messages

The **Status Messages** tab displays messages showing the status of various steps. Selecting a step name drills into the step's messages, which could include user authorization warnings or step processing errors.

Log

The **Log** tab displays the log for the selected job.

SAS® and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration. Other brand and product names are trademarks of their respective companies. Copyright © 2015, SAS Institute Inc., Cary, NC, USA. All Rights Reserved. January 2023 2.2-P3:evjobmtr

THE POWER TO KNOW.