

Scheduler Guide

Axiom

Version 2021.2

AXIOM

Microsoft, Excel, Windows, SQL Server, Azure, and Power BI are trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries.

This document is Syntellis Performance Solutions, LLC Confidential Information. This document may not be distributed, copied, photocopied, reproduced, translated, or reduced to any electronic medium or machine-readable format without the express written consent of Syntellis Performance Solutions, LLC.

Copyright © 2021 Syntellis Performance Solutions, LLC. All rights reserved. Updated: 7/21/2021

Syntellis
10 S. Wacker Dr., Suite 3375
Chicago, Illinois 60606
847.441.0022
www.syntellis.com

Table of Contents

| | |
|---|----|
| Chapter 1: Introduction | 1 |
| Chapter 2: Scheduler Overview | 3 |
| About Scheduler | 3 |
| The Scheduler dialog | 6 |
| Chapter 3: Scheduler Setup | 9 |
| Enabling Scheduler services | 9 |
| Configuring Scheduler system jobs | 10 |
| Setting up email delivery for Axiom | 12 |
| Chapter 4: Scheduler Job Setup | 20 |
| Managing Scheduler jobs and tasks | 20 |
| Defining scheduling rules for a job | 23 |
| Setting up notifications for jobs | 28 |
| Creating event handlers for a job | 31 |
| Using job variables | 33 |
| Processing tasks in parallel | 36 |
| Using iterative task processing | 38 |
| Conditionally processing tasks in a job | 42 |
| Run another Scheduler job from within a Scheduler job | 44 |
| Running a job | 46 |
| Using RunEvent to execute a Scheduler job | 46 |
| Job properties | 52 |
| Chapter 5: Scheduler Administration | 60 |
| Managing scheduled jobs | 60 |
| Viewing job results | 62 |
| Managing Scheduler servers | 63 |
| Managing event handlers | 66 |
| Managing remote data connections | 67 |
| Chapter 6: Scheduler Task Reference | 70 |
| Task Control properties | 71 |
| Active Directory Import task | 73 |
| Collect Worksheets task | 77 |
| Copy On Demand Plan Files task | 79 |

| | |
|---|------------|
| Create Plan Files task | 86 |
| Echo task | 88 |
| Execute Command Adapter task | 89 |
| Execute SQL Command task | 90 |
| Export ETL Package task | 91 |
| File Processing task | 91 |
| Import ETL Package task | 96 |
| Process Plan Files task | 97 |
| Process Document List task | 107 |
| Process Template List task | 112 |
| Purge System Data task | 115 |
| Raise Event task | 116 |
| Run Scheduler Job task | 117 |
| SMTP Message Delivery task | 119 |
| Start Process task | 120 |
| Web Report Processing task | 121 |
| Scheduler tasks for database maintenance | 133 |
| Chapter 7: Web Scheduler | 134 |
| Managing the job schedule in the Web Client | 135 |
| Viewing job results in the Web Client | 138 |
| Running a job manually in the Web Client | 141 |
| Viewing jobs and event handlers in the Web Client | 142 |
| Index | 147 |



Introduction

Axiom's Scheduler is used to schedule tasks for automated execution and perform those tasks using server-side processing. Various tasks can be scheduled, such as processing plan files, file processing, importing data, and more.

► Intended audience

This guide is intended for administrators and other power users who are responsible for scheduling tasks for automated execution.

► What is covered in this guide?

This guide covers the following aspects of Scheduler setup and management:

- Performing initial Scheduler setup
- Creating jobs and tasks and scheduling them for execution
- Using advanced Scheduler features such as variables and event handlers
- Reference for all job and task properties

► What is not covered in this guide?

The following related topics are not covered in this guide:

- Security setup for Scheduler. For more information, see the *Security Guide*.
- Explanation of certain features that can be scheduled. For example, if you want to schedule a data import, this guide assumes that you already know what a data import is and that the import package has already been created. This guide only explains how to set up that import for execution in Scheduler.

All documentation for Axiom can also be accessed using the Axiom Help Files.

► Axiom Client versions

Most of this guide discusses Scheduler functionality that is available in the Axiom Desktop Client (Excel Client and Windows Client). Screenshots of features may show either the Excel Client or the Windows Client. The Axiom functionality is virtually identical in both environments.

Scheduler also has limited functionality available in the Web Client. Any web-specific screenshots and features are clearly noted.

Scheduler Overview

Using Scheduler, you can schedule certain Axiom tasks to be processed on a Scheduler server at a specific date and time. For example, you can schedule plan file processing or data imports.

Processing tasks using Scheduler has advantages over manual processing, such as:

- Leverages the server's processing power and frees up your computer's resources.
- Enables recurring scheduling of ongoing tasks.
- Allows tasks to be scheduled during "off hours," during periods of low network and system activity.
- Allows tasks to be performed in batch, including enforcing task dependencies.

Scheduler processes tasks using jobs. Each job is a scheduled unit that can contain one or more tasks. The tasks in a job can be processed sequentially or concurrently as appropriate.

Only system administrators and users with the **Scheduled Jobs User** security permission can access Scheduler.

Most Scheduler setup activities can only be performed in the Desktop Client (Excel or Windows Client). Therefore, the Desktop Client Scheduler is the primary focus of this document. However, some job management activities can be performed in the Web Client, such as monitoring the job schedule, viewing job results, and running jobs manually on demand. For more information, see [Web Scheduler](#).

About Scheduler

This section contains conceptual information about the Scheduler feature in Axiom.

► Scheduler jobs and tasks

The primary unit of Scheduler processing is a *job*. Each Scheduler job can contain one or more *tasks* to be performed as part of that job.

Each Scheduler job defines the following basic properties:

- The tasks to perform for the job and the properties of those tasks
- The schedule of the job, including recurrence (if any)
- The priority of the job

- The notification options for the job

The tasks define the actual activities to be performed by the job, such as importing data or processing plan files. Some Scheduler tasks correspond to existing features that can also be processed manually (such as Process Plan Files), while other tasks are Scheduler-specific and can only be processed via Scheduler. Each task has a unique set of options that are specific to that task and to the activity to be performed. For more information on the available task types, see [Scheduler Task Reference](#).

The tasks in a job can be processed sequentially or concurrently as appropriate. Tasks can be dependent on other tasks in the job as needed—for example, you can configure a job so that if a task fails, the job stops and does not process the next task. Tasks can also be processed iteratively, to perform the same task repeatedly over a defined set of values.

The Scheduler jobs in your system fall into the following basic categories:

- **Client-created:** You can create Scheduler jobs as needed to perform tasks in your system.
- **System jobs:** Axiom provides a set of [system jobs](#) to perform necessary system tasks.
- **Product-controlled:** When a product is installed, it may include one or more Scheduler jobs to support the use of that product. Generally speaking, these jobs should not be changed unless the product documentation says customization is allowed, or as advised by Axiom Support.

► How Scheduler jobs are run

Once a Scheduler job has been created, it can be run using any of the following options:

- The job can be scheduled for execution at a future date and time using a scheduling rule. Scheduling rules can be one-time only, or recurring.
- The job can be run "one time" manually as needed through Scheduler.
- The job can be triggered for execution using an event handler. This allows Scheduler jobs to be triggered in various ways, such as by clicking a button in an Axiom form.

Scheduler jobs are processed by one or more servers running the Scheduler service. For Axiom Cloud systems, the Scheduler service is part of your cloud system and managed by Axiom Support. For on-premise systems, the Scheduler service is installed on one or more servers in your environment. The Scheduler service polls the Axiom Application Server periodically to check for any jobs that are ready to be run. Eligible jobs are then executed on the server, based on their [processing priority](#).

When a job is executed by Scheduler, it is run using a particular user identity. In order for a job to be executed successfully, the user must be an active user defined in Axiom security, and the user must have the appropriate security permissions to perform the tasks in the job. The user identity for a job is determined as follows:

- If a job is a system job, then it is run as the system-managed identity of **System** instead of a user identity.
- If a job is run by using **Run Now**, then it is run as the user who placed the job on the schedule.
- If a job is run by an active scheduling rule, then it is run as the *job owner*. The job owner is the user who last saved the job.

- If a job is run via an event handler, then the job may be run as either the job owner, or the job requester (the user who raised the event).

► System jobs

System jobs are automatically created by Axiom to support necessary system functionality. Some system jobs are created as part of the initial installation and are intended to run on an ongoing basis, while other system jobs are created on-demand in response to system events. Only administrators can edit these system jobs.

System jobs have two defining characteristics:

- System jobs are run using the system-managed identity of **System** instead of a user identity. The System identity has full rights to the system as necessary to perform system tasks.
- System jobs are run by the default System Scheduler service. For on-premise systems, this service is created and started automatically on the Axiom Application Server, and does not require a separate installation. This service is exclusively for running system jobs.

Axiom Cloud systems may or may not have a separate System Scheduler service, depending on the system configuration (as determined by Axiom Support). If your cloud system does not have a System Scheduler service, then your system jobs are run using the available Scheduler services for the cloud system.

In the Scheduler dialog (Desktop Client), the System Scheduler service is listed on the **Servers** tab using the following naming convention: **<ServerName>-System**.

If necessary, a product-controlled or client-created job can be flagged as a system job, so that it can be run using the System identity instead of a user identity. To designate a job as a system job, enable **Mark as System Job** in the **General** job properties. The following rules apply to manually-created system jobs:

- Only system administrators can designate a job as a system job.
- The job cannot contain any tasks that are designated as “non-system” tasks. Non-system tasks are any tasks that might involve spreadsheet processing, such as Process Plan Files.

► Processing priority for scheduled jobs

Once a job reaches its start time, it is eligible to be processed by Scheduler and joins the processing queue. For scheduled jobs, the start time is based on the scheduling rule that placed it on the schedule. For other jobs, the start time is the time that the job was placed on the schedule using **Run Once** or triggered by an event handler.

Each Scheduler service has a configured number of *threads* that are used to process jobs. As a Scheduler thread becomes available, it takes the next job in the processing queue. The priority of jobs in the processing queue is determined by the combination of the job's priority category, and its **Priority Elevation** setting.

Each job has a priority category, based on how the job execution was initiated. The priority categories are as follows:

1. **Manual:** The job was executed manually.
2. **Event Handler:** The job was executed by a Scheduler event handler.
3. **Scheduled Job:** The scheduled instance of the job results from an active scheduling rule.
4. **Subordinate Job:** The job was generated as a subordinate job, from a currently executing job.

Manual jobs are highest priority and are processed first, and subordinate jobs are lowest priority and are processed last. Within each category, jobs are processed according to their **Priority Elevation** setting.

For example, imagine that Scheduler has 2 available threads and the following jobs are eligible to be processed:

| Job | Priority Category | Priority Elevation |
|-----|-------------------|--------------------|
| A | Manual | Default |
| B | Event Handler | Default |
| C | Scheduled | Default |
| D | Scheduled | Elevated |

- Scheduler will execute jobs A and B first, because those are the highest priority jobs based on their priority category.
- When the next thread becomes available, Scheduler will execute job D. Although job C may have entered the queue first, and the two jobs have the same priority category, job D's priority elevation is set to **Elevated** so it takes precedence within the category. If instead both jobs were set to **Default**, then job C would be executed first if it entered the queue before job D.
- When the next thread becomes available, Scheduler will execute job C.

NOTE: If a job's **Priority Elevation** is set to **Interrupt**, then it is run as soon as it is eligible, regardless of its priority category and regardless of whether any Scheduler threads are currently available to process the job. If no Scheduler threads are available, a new one is created to process the job, even if this temporarily exceeds the number of configured threads for the server.

The Scheduler dialog

The **Scheduler** dialog is used to create and manage Scheduler jobs.

To access Scheduler:


- On the **Axiom** tab, in the **Administration** group, click **Manage > Scheduler**.

NOTE: In systems with installed products, this feature may be located on the **Admin** tab. In the **System Management** group, click **Scheduler**.

The top section of the **Scheduler** dialog contains a ribbon-style toolbar with two tabs: **Job** and **Service**.

- On the **Job** tab, you can create, run, and edit jobs.
- On the **Service** tab, you can manage scheduled jobs, view job results, and perform other Scheduler management activities.

As you perform actions on the **Job** and **Service** tabs, additional tabs are opened in the navigation pane of the dialog. For example, clicking the **Scheduled Jobs** button on the **Service** tab opens the **Scheduled Jobs** tab in the navigation pane. You can move between any open tab in the navigation pane, regardless of which tab is selected in the ribbon. The ribbon updates to show the related commands for the selected item.



Example Job tab

| ID | Job | User | Status | Server | Priority | Start Time | Due In |
|---------|----------------------------|--------|---------|--------|---------------|-----------------|---------------|
| 4559790 | System.ProcessNotification | System | Pending | | Scheduled Job | 7/11/2018 17:01 | 44.42 minutes |
| 4559766 | System.SystemDataPurge | System | Pending | | Scheduled Job | 7/12/2018 01:00 | 8.72 hours |
| 4559769 | System.IndexMaintenance | System | Pending | | Scheduled Job | 7/12/2018 05:15 | 12.97 hours |
| 4559792 | Process Plan Files | admin | Pending | | Scheduled Job | 7/12/2018 16:05 | 23.81 hours |

Example Service tab

When you right-click a tab in the dialog's navigation pane, you can close or save items as follows:

- For all items, you can **Close**, **Close All**, or **Close All But This**.
- For jobs, you can **Save** or **Save As**. Selecting **Save As** allows you to save a copy of the job to the **Scheduler Jobs Library** in the Axiom file system.

The **Scheduler Jobs Library** is also accessible via Axiom Explorer.

Scheduler Setup

If this is a new Axiom installation, certain setup activities need to be performed to enable Scheduler functionality for each system.

Enabling Scheduler services

All Axiom systems must have one or more Scheduler services available to execute Scheduler jobs. These Scheduler services must be explicitly enabled within the Scheduler dialog before they are available for use.

- For on-premise systems, these are the Scheduler services that you manually installed as part of the Axiom server installation.
- For Axiom Cloud systems, these Scheduler services were created for you when your Axiom Cloud system was created.

To enable Scheduler services:

1. In the **Scheduler** dialog, on the **Service** tab, click **Servers**.

You should see one or more servers listed in this tab, representing your manually installed services or the services created for your Axiom Cloud system. These are the servers that you must enable for processing.

You may also see the following items in the **Servers** tab:

- **AppServerName-System**: For on-premise systems, this is the system Scheduler service that was installed automatically as part of your Axiom Application Server installation. This service should already be enabled. It does not require any special setup.
 - **ServerName-CIS**: For Axiom Cloud systems, this is the Cloud Integration Service that was manually installed to facilitate connections between your local environment and your cloud system. For more information on this service, see the *Axiom Cloud Technical Guide*.
2. Select each server in the list. In the **Configuration details** section at the bottom of the dialog, select the **Schedule Processing Enabled** check box, and then click **Update**.

NOTE: The **Configuration details** section also contains several configuration settings for the Scheduler service. We recommend leaving the default settings unless you are instructed to change them by Axiom Support. For more information, see [Managing Scheduler servers](#).

These Scheduler servers are now available to process Scheduler jobs.

Configuring Scheduler system jobs

Scheduler includes several built-in system jobs that should be configured as part of your system setup activities. These jobs are created automatically by the system Scheduler service when it is started. These jobs control email notifications for Scheduler and other system processes, and perform periodic database cleanup.

Only administrators can edit these system jobs.

To configure the Scheduler system jobs:

1. In the **Scheduler** dialog, on the **Service** tab, click **Scheduled Jobs**. You should see the following scheduled jobs where the name starts with **System**:
 - **System.ProcessNotification**
 - **System.SystemDataPurge**
 - **System.IndexMaintenance**
2. Double-click a job to open it for editing. Once you have modified the job as needed, click **Save**.
 - In the **Tasks** section, select the task name to access the task settings, and then edit the **Task Details** as necessary. See the following section for more details on the specific task settings.
 - In the **Notifications** section, configure the email notifications for the job as desired. By default, the jobs are configured to send notifications on error, however, you must specify a valid email address to receive these notifications (such as the email address of a system administrator).
3. On the **Service** tab, click **Event Handlers** and then locate the **System.SMTPMessageDelivery** event in the list. Double-click the event to open the associated job (with the same name), then configure the task settings. For more information, see [Setting up email delivery for Axiom](#).

The changes will apply the next time the job is executed.

TIP: Administrators can also access these jobs in Axiom Explorer, in `\Axiom\Axiom System\Scheduler Working Folder`.

If one of these jobs is removed from the schedule or deleted, you can restore it by right-clicking any system job or event handler and selecting **Create system jobs and event handlers**. Any unscheduled or deleted system jobs will be restored.

► System job overview

The following list provides a brief description of the purpose of each job and its task settings.

SMTPMessageDelivery

| | |
|---------------|---|
| Description | This job enables email notifications for Scheduler jobs and other system processes, such as process management notifications. Email notifications are stored in the database until the job picks them up and sends them. |
| Configuration | This task is not operational until you specify the name of a valid SMTP server in your environment. Once you have completed the SMTP server settings, make sure to disable Test Mode so that email notifications will be sent. |
| Schedule | This job is not a scheduled job. Instead, it uses an event handler so that the job is triggered for execution whenever there are email messages ready to be sent. |

ProcessNotification

| | |
|---------------|--|
| Description | This job enables reminder notifications for process management. These are notifications to remind users of upcoming due dates and overdue tasks. |
| Configuration | This task does not have any editable settings. |
| Schedule | By default, this job is set to execute once per hour. You should not change the frequency of this job. |

SystemDataPurge

| | |
|---------------|---|
| Description | This job purges old data from the system, such as old Scheduler job results, old email notifications, old temp table data, and old audit data. This task should be run regularly to help keep old and unnecessary data from impacting system performance. |
| Configuration | This task is configured to delete items that are older than a specified number of days (by default, 15 days). You can change the default settings for this task if desired. |
| Schedule | By default, this job is set to run every ten minutes from midnight to 4:00 AM. The task is operational as soon as it is created by the Scheduler service, using the default settings. |

IndexMaintenance

| | |
|---------------|--|
| Description | This job maintains the indexes on your Axiom databases (system database and audit database). It reports the current index fragmentation and then reorganizes or rebuilds the indexes as needed, depending on the current level of fragmentation. It also updates table statistics. |
| Configuration | Any changes made to the SQL statements in this system job should only be made with the guidance of Axiom Support and/or your SQL database administrator. |
| Schedule | By default, this task is set to run once per day, at 5:15 AM. The task is operational as soon as it is created by the Scheduler service, using the default settings. |

► Email notifications for system jobs

By default, most system jobs are configured to send notifications on error, however, you must edit the job to specify a valid email address for the notifications (such as the email address of a system administrator). In the **Notifications** section for the job, edit the email notification settings as desired. The default setting of `{CurrentUser.EmailAddress}` cannot be used because the job is run by the system and therefore does not resolve to a user email address.

NOTE: It is recommended to leave the `SMTPMessageDelivery` job at the default notification behavior of `None`. If this job experiences an issue attempting to send email, it likely will be unable to send you an email notification about this error.

For more information on notification settings, see [Setting up notifications for jobs](#).

Setting up email delivery for Axiom

Various features in Axiom can be configured to send email. In order to enable email delivery for Axiom, the **System.SMTPMessageDelivery** event handler in Scheduler must be set up to use a mail server in your environment.

Whenever an Axiom feature adds an email to the delivery queue, the SMTP event handler triggers a Scheduler job to send the email. Typical emails include Scheduler job status notifications, report delivery using file processing, and task notifications for process management.

NOTE: These setup steps only apply to on-premise systems. If you have an Axiom Cloud system, email functionality will be configured for you.

► Prerequisites

In order to set up the SMTP event handler, you must know the following information about a mail server in your environment. Consult with your Information Technology department to obtain this information:

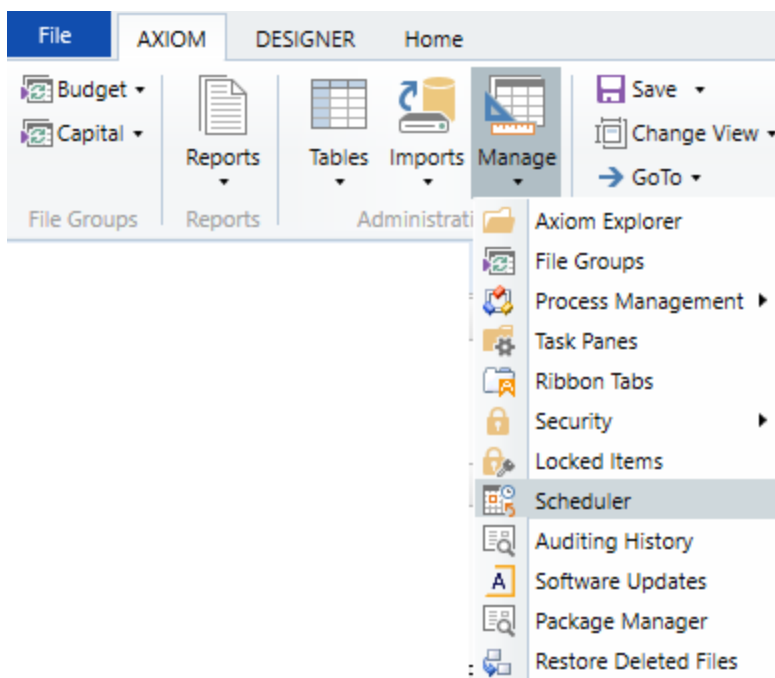
- The name of the SMTP email server
- The port number to use to send email from the server (by default, port 587 is used)
- An account name and password to access the server

► Configuring the SMTP event handler

The SMTP event handler can only be edited in the Desktop Client, and only by an administrator.

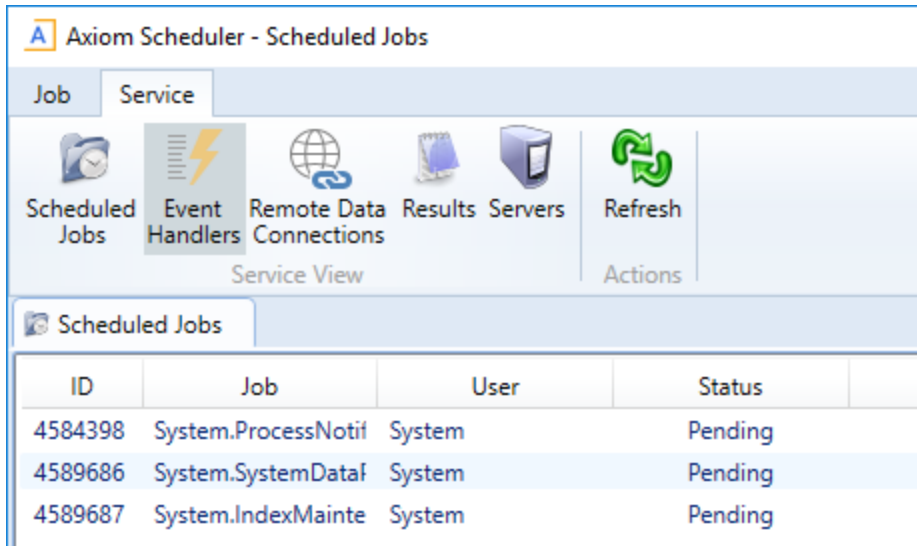
To configure the SMTP event handler:

1. On the **Axiom** tab, in the **Administration** group, click **Manage > Scheduler**.

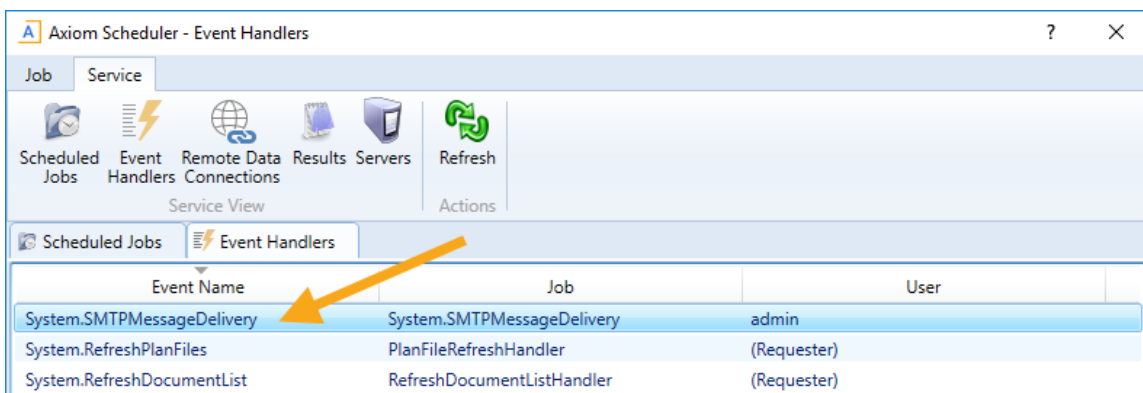


NOTE: In systems with installed products, this feature may be located on the **Admin** tab. In the **System Management** group, click **Scheduler**.


2. In the **Scheduler** dialog, on the **Service** tab, click **Event Handlers**.



3. In the list of event handlers, double-click **System.SMTPMessageDelivery** to open the associated job.



4. In the **Tasks** section of the job, select the **SmtpMessageDeliveryTask**.



5. In the **Task Details**, complete the **SMTP Server Information**:

| Item | Description |
|--------------------------------|--|
| Server name | The server name of the SMTP email server. |
| Port number | The port number for the SMTP email server. By default, the port is 587, but you can specify a different port number if needed. |
| Server requires authentication | Select this check box if the SMTP email server requires authentication. If selected, type a Username and Password . |

6. If you are ready to start sending email, clear the **Test Mode** check box.

When the SMTP task is in test mode, it processes email messages in the queue but does not actually deliver them via email. The SMTP task is set to test mode by default, so that it does not attempt to send email before it is properly configured. Once you have configured the task with the necessary mail server information, you can disable test mode to start sending email.


IMPORTANT: Once the job is saved with **Test Mode** disabled, email will start being sent from Axiom. Make sure that you do not have any active features left over from system testing that might start generating unwanted email notifications.

7. Click **Save**.

► Testing email delivery


To test that your SMTP event handler is configured correctly, you can generate an email using the **E-mail** feature.


1. Open a spreadsheet report in the Desktop Client.
2. On the **Axiom** tab, in the **File Output** group, select **E-mail**.



NOTE: In systems with installed products, this feature may be located on the **Main** tab—either directly on the ribbon or under **Publish**.

3. In the **E-mail Active Workbook** dialog, select the following options, then click **OK**:
 - For **Send As**, select **Document Link**.
 - For **Send using**, select **Axiom Mail Service**.

 E-mail Active Workbook ? ×

 Send a snapshot of or a document link to Home.xlsx.

Send As: ☐ Snapshot ☒ Document Link

Send using: ☐ Mail ☒ Axiom Mail Service


Snapshot Options:

Send file as: XLSX - Microsoft Excel Worksheet (.xlsx) ▾

Include: ☒ Active Worksheet Only
☐ Entire Workbook

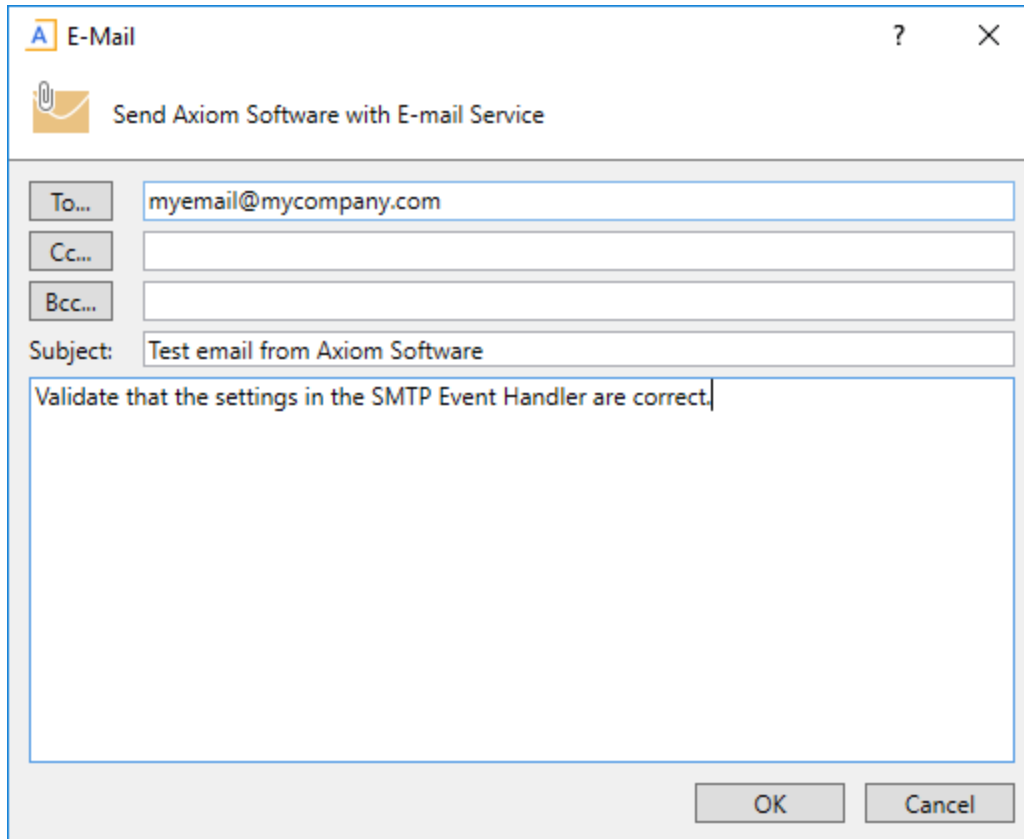
Formulas: ☒ Convert All Formulas
☐ Retain Excel Native Formulas

Document Link Options:

Sheet Filter: 


Cell Address:

4. In the Email dialog, enter your email address, and complete the subject and body text to indicate that you are sending a test email. Click **OK** to send the email.



You should receive the test email within a minute or two, depending on Scheduler availability. If you do not receive the test email, you can review the Scheduler job results to determine what happened:

1. Re-open the Scheduler dialog, and re-open the **System.SMTPMessageDelivery** event handler (as [previously described](#)).
2. Select **Results**, then review the list of results to locate the execution that should have delivered your email. (The **Started** time should help you determine the right set of results, if several results are present.) Expand the result to view the details. If the execution failed, check the error message to see what needs to be changed in the SMTP event handler.



- Update the SMTP event handler job as needed, and then repeat the test. If the test still fails, or if the error details do not seem to be related to the SMTP server configuration, contact Axiom Support for further assistance.

Scheduler Job Setup

To perform Axiom tasks using Scheduler, you must create jobs. Each job can execute one or more tasks. This section discusses how to set up jobs, including how to schedule jobs for future execution and how to be notified when a job has been completed.

Managing Scheduler jobs and tasks

Using the **Axiom Scheduler** dialog, administrators can create and edit Scheduler jobs. To access this dialog:

- On the **Axiom** tab, in the **Administration** group, click **Manage > Scheduler**.

NOTE: In systems with installed products, this feature may be located on the **Admin** tab. In the **System Management** group, click **Scheduler**.

This section discusses how to create, edit, and delete jobs and tasks, not how to manage the Scheduler queue once jobs have been placed on the schedule. If you need to stop or reschedule a scheduled job, see [Managing scheduled jobs](#).


Scheduler jobs are saved as XML files and are stored in the Axiom file system at `\Axiom\Scheduler Jobs Library`.

► Creating a Scheduler job

You can create a new Scheduler job to perform one or more tasks.

To create a new job:

1. In the **Scheduler** dialog, on the **Job** tab, click **New**.



A new tab appears in the navigation pane, labeled **New Job**. The left-hand side of the job lists sections for which you can define various job settings. When you click a section name, the settings for that section display in the right-hand side of the job.

2. In the **General** section, define general job settings as desired.

For detailed information on the available settings for a job, see [Job properties](#).

3. In the **Scheduling Rules** section, specify scheduling details for the job.

You can schedule the job for future execution, for one time or on a recurring basis.

NOTE: If you are always going to run the job manually, and do not need to schedule it for future execution, then you do not need to define scheduling rules.

For more information, see [Defining scheduling rules for a job](#).

4. In the **Notification** section, specify email notification options for the job.

You can send email notifications every time the job completes, or only when the job experiences errors. By default, the job is configured to notify on completion.

For more information, see [Setting up notifications for jobs](#).

5. In the **Tasks** section, add one or more tasks to the job.
 - a. On the **Job** tab of the ribbon, in the **Tasks** group, click **Add**. This brings up a list of available tasks. Select the task that you want to add.

The task is added to the **Tasks** section, and the settings for the task display in the right-hand side of the job.

- b. Complete the settings for the task as desired.

The **Task Control** section of the task contains standard task settings, and the **Task Details** section contains settings unique to the task type. For more information, see [Task Control properties](#).

If a required setting is not completed, the setting is highlighted in red and error text appears in the bottom of the dialog. Make sure to complete all required settings for the task before saving.

Repeat this process until you have added all desired tasks to the job. Tasks are processed in the order listed. If you need to change task order, select a task and then click **Move Up** or **Move Down**.

6. In the **Job** tab of the ribbon, click **Save**.
7. At the bottom of the **Save As** dialog, in the **File name** box, type a name for the job, and then click **OK**.

The job is saved as an XML file in the Scheduler Jobs Library.

If the job was saved with an active scheduling rule, Axiom determines the next scheduled date of execution and schedules the job. You can view the job in the **Scheduled Jobs** list (on the **Service** tab of the ribbon, click **Scheduled Jobs**).

Advanced job settings

This procedure covers the basic steps of creating a job. Jobs also support the following advanced options:

- **Event handlers:** You can create event handlers for the purposes of running the job using the RunEvent function. This allows users to trigger job execution from within an Axiom file.
- **Job variables:** You can create job variables and then use those variables within certain job settings. You can then dynamically pass in values for those variables when using the RunEvent function to execute the job.

For more information, see [Creating event handlers for a job](#), [Using job variables](#), and [Using RunEvent to execute a Scheduler job](#).

► Editing a job

You can edit a job at any time to change job settings, add or remove tasks, change scheduling rules, or change notification options.

This section describes the general process of opening a job for editing. For more details on the impacts of editing scheduling rules, see [Defining scheduling rules for a job](#).

To edit a job:

1. In the **Scheduler** dialog, in the **Job** tab, click **Open**.

The **Axiom Explorer** dialog opens, showing the Scheduler Jobs Library only.

2. Select the job and then click **Open**.

The job opens in the **Scheduler** dialog. Make sure the job is the active tab in the navigation pane (the most recently opened tab is the active tab by default).

3. Edit the job and task properties as desired.

For detailed information on the available settings for a job, see [Job properties](#). For detailed information on task settings, see [Task Control properties](#).

4. In the **Job** tab of the ribbon, click **Save**.

► Deleting a job

Deleting a job removes any scheduled executions of the job from the scheduled jobs list.

To delete a job:

1. In the **Scheduler** dialog, in the **Job** tab, click **Open**.

The **Axiom Explorer** dialog opens, showing the Scheduler Jobs Library only.

2. Navigate to the job that you want to delete, then right-click the job and select **Delete**.

You can also delete Scheduler jobs from Axiom Explorer or the Explorer task pane.

Defining scheduling rules for a job

Once a job has been created, you can run it on demand, or you can schedule it for future execution. Jobs can be scheduled to be run one time, or on a recurring basis. To schedule a job, you define scheduling rules for the job.

You can add, edit, and remove the scheduling rules for a job at any time using the **Scheduling Rules** section of the job properties. You can also flag a rule as active or inactive. If a job has no scheduling rules, or if all of its scheduling rules are inactive, then it will not be run unless it is run manually by a user.

If a job is saved with an active scheduling rule, then Axiom determines the next scheduled instance of the job and places it in the scheduled jobs list. Once that instance has been processed, the next scheduled instance is determined and scheduled, and so on. Each time the job is run using an active scheduling rule, it is run as the current job owner (unless the job is a system job, in which case it is run as the Scheduler Service System identity).

If a job has multiple active scheduling rules, Axiom evaluates all of the rules and schedules a single instance of the job, for the earliest time allowed by the rules. Multiple scheduling rules do not result in multiple scheduled instances of the job.

NOTE: If a time zone is listed on the Scheduling Rules section of the job, then the defined rules will be evaluated in the context of that listed time zone. Otherwise, scheduling rules are evaluated in the context of the local time zone for the Scheduler Server. If necessary, the system configuration setting **SchedulingBehaviorTimezone** can be used to specify a particular time zone for evaluating scheduling rules.

► Adding a Scheduling rule

You can add a scheduling rule to a job to schedule it for future execution, either one time or on a recurring basis.

If you only plan to run the job manually on demand, then you do not need to create a scheduling rule.

To add a scheduling rule to a job:

1. In the **Scheduler** dialog, open a job to edit or create a new job.
2. In the left-hand side of the job, select **Scheduling Rules**.

By default, this area is empty. You must add a rule in order to define scheduling for the job.

3. On the **Job** tab of the Scheduler ribbon, in the **Scheduling Rules** section, click **Add**.

A new row appears in the right-hand side of the job. By default, the new row is active, but does not have start / end dates or any specific recurrence settings.

4. Complete the following settings within the row as needed:

| Item | Description |
|--------|--|
| Active | <p>If you want the job to be placed on the schedule as soon as you save the job with the new scheduling rule, then you should leave this option checked.</p> <p>However, if you just want to save your schedule settings but you are not ready to begin scheduling the job, then you can clear the Active check box for the rule. The job will not be scheduled until it is saved with an active scheduling rule.</p> |

| Item | Description |
|--------------------------|--|
| Starting On Ending On | <p>Optional. These dates specify the time frame for the scheduling rule. The starting date defines the earliest point in time that the job can be scheduled, and the ending date defines the latest point in time that the job can be scheduled.</p> <p>If these dates are not defined (left blank), then the job will be perpetually scheduled according to the rule settings, as long as the rule is active.</p> <p>If you want to schedule a one-time job, then set the starting / ending dates to the same date and time.</p> <p>NOTE: Your system locale determines the format of dates.</p> |
| Day of Week | <p>Specify the day(s) of the week that you want the job to be run:</p> <ul style="list-style-type: none"> • * (Default): The job will be run on all days within the start / end range. • 0-6: The job will be run on the specified day or days, where 0 is Sunday and 6 is Saturday. Use a comma or a hyphen to separate multiple days (hyphen if the days are contiguous, commas if not). <p>For example, you can enter 1, 3, 5 for Monday, Wednesday, and Friday, or enter 1-5 for Monday through Friday.</p> |
| Hours | <p>Specify the time of day (hours) that you want the job to be run, in relation to the specified days:</p> <ul style="list-style-type: none"> • * (Default): The job will be run on all hours. • 0-23: The job will be run on the specified hour or hours, where 0 is midnight and 23 is 11:00 PM. Use a comma or a hyphen to separate multiple hours (hyphen if the hours are contiguous, commas if not). <p>For example, you can enter 0, 12 to run at midnight and noon, or enter 0-12 to run every hour from midnight to noon.</p> |

| Item | Description |
|---------|---|
| Minutes | <p>Specify the time of day (minutes) that you want the job to be run, in relation to the specified hours:</p> <ul style="list-style-type: none"> • * (Default): The job will be run on all minutes (essentially the job is run continuously, once per minute). • 0-59: The job will be run on the specified minute or minutes of the hour, where 0 is the first minute of the hour and 59 is the last minute of the hour. Use a comma or a hyphen to separate multiple minutes (hyphen if the hours are contiguous, commas if not). <p>For example, you can enter 0, 30 to run at the top of the hour and the half hour, or enter 0-30 to run every minute from the top of the hour to the half hour.</p> <p>NOTE: If you specify an hour, then in most cases you should also specify a minute (such as 0 to run the job at the top of the specified hour). If you enter an hour but leave the minutes at the default asterisk, then the job will run every minute in that hour.</p> |

If the Active check box for the rule is selected when the job is saved, then Axiom will calculate the date and time of the first scheduled execution and will place the job on the schedule.

► Editing a scheduling rule

You can edit a scheduling rule at any time, to toggle between active and inactive, and to change the start / end dates and recurrence settings.

NOTES:

- If a pending instance of this job is currently on the schedule, and you edit the scheduling rule, the pending instance will be updated to match the new schedule.
- If you inactivate a scheduling rule, any currently scheduled instances of the job will be automatically removed from the schedule.

To edit a scheduling rule:

1. In the **Scheduler** dialog, open a job to edit or create a new one.
2. In the left-hand pane of the job, select **Scheduling Rules**.
The defined rules display in the right-hand pane of the job.
3. Make any desired changes directly within the scheduling rules grid.

► Deleting a scheduling rule

You can delete a scheduling rule at any time. If a job has no active scheduling rules, it will not be processed unless it is manually run.

To delete a scheduling rule:

1. In the **Scheduler** dialog, open a job to edit.
2. Select the **Scheduling Rules** section of the job, and then select the rule that you want to delete.
3. On the **Job** tab of the Scheduler ribbon, in the **Scheduling Rules** group, click **Remove Selected**.

Alternatively, if you want to delete all scheduling rules for the job, click **Clear All**.

Any jobs in the scheduled jobs queue that were related to the deleted rule(s) are also deleted.

► Scheduling rule examples

The following are some example schedules and the rules used to achieve them:

| Schedule | Start/End | Day of Week | Hours | Minutes |
|--|--|-------------|-------|------------|
| Weekdays at 11:00 PM | <optional> | 1,2,3,4,5 | 23 | 0 |
| Every 15 minutes | <optional> | * | * | 0,15,30,45 |
| Mondays at 11:30 PM | <optional> | 1 | 23 | 30 |
| One time (6/30/2021) at 1:30 PM (Option 1) | Start: 06/30/2021 00:00 End: 07/01/2021 00:00 | * | 13 | 30 |
| One time (6/30/2021) at 1:30 PM (Option 2) | Start: 06/30/2021 13:30 End: 06/30/2021 13:30 | * | * | * |
| Every Wednesday in July at noon | Start: 07/01/2021 00:00 End: 08/01/2021 00:00 | 3 | 12 | 0 |
| Continuous | <optional> | * | * | * |

To schedule a job to execute monthly, create twelve active scheduling rules, one for each month. This is necessary because scheduling rules do not have a property for day of month, so it is not possible to use a single scheduling rule to create a monthly schedule. In the following example, the job will be executed on the first day of each month, at 3:30 AM:

| Active | Starting On | Ending On | Day Of Week | Hours | Minutes |
|-------------------------------------|----------------|----------------|-------------|-------|---------|
| <input checked="" type="checkbox"/> | 1/1/2021 00:00 | 1/2/2021 00:00 | * | 3 | 30 |
| <input checked="" type="checkbox"/> | 2/1/2021 00:00 | 2/2/2021 00:00 | * | 3 | 30 |
| <input checked="" type="checkbox"/> | 3/1/2021 00:00 | 3/2/2021 00:00 | * | 3 | 30 |
| <input checked="" type="checkbox"/> | 4/1/2021 00:00 | 4/2/2021 00:00 | * | 3 | 30 |

Example scheduling rules to execute a job monthly

When you save the job, the rules will be evaluated and the first scheduled execution will be placed on the schedule—in this example, the January 1 execution. Once that scheduled execution is complete, the rules will be evaluated again, which will cause the next scheduled execution (Feb 1) to be placed on the schedule, and so on.

Setting up notifications for jobs

Scheduler can be configured to send an email notification when a job completes, or when a job has errors. In order for an email notification to be sent for a particular job, the following must be set up:

- The job must be configured to send a notification on completion or error. The notification settings must include valid To and From email addresses (or use system variables).
- The System.SMTPMessageDelivery system job must be configured with a valid SMTP server for your environment. For more information, see [Configuring Scheduler system jobs](#).

When a job creates an email notification, the notification is first saved to the database. When notifications are detected in the database, the System.SMTPMessageDelivery system job is triggered to deliver the notifications.

NOTES:

- By default, all new Scheduler jobs are configured to send an email notification on completion, to the user who created the job. You only need to edit the notification settings if you want the job to use different notification behavior.
- Currently, it is not possible to configure a Scheduler job to send notifications within the application only, instead of by email. However, when a job is run manually, the user who ran the job may receive an in-application notification of the job status in addition to any configured email notifications. See [Application notifications for Scheduler jobs that are run manually](#).

To configure a job to send email notifications:

1. In the **Scheduler** dialog, open a job to edit or create a new job.
2. In the left-hand side of the job, select **Notification**.
3. In the **Job Notification Level** section, select one of the following:
 - **Send all email notifications:** (Default) An email notification is always sent when the job is executed, regardless of the job status (success, failure, aborted, etc.).

- **Send email notification only when the job has errors:** An email notification is only sent if the job experiences errors. If the job completes successfully with no errors, no email notification is sent.
- **None:** No email notifications are sent for this job. The only way to check the status of the job execution is to view the job history.
- **Send email notification to different email addresses when the job has errors or succeeds:** This option works the same way as **Send all email notifications**, except that a separate email address can be specified to receive the error notifications.

4. In the **Notification Message Content** section, complete the following for the notification email:

| Item | Description |
|------|--|
| To | <p>The email address(es) to receive the notification email. Separate multiple addresses with a semicolon. For example:</p> <ul style="list-style-type: none"> • To send the email to two recipients, enter the addresses such as: jdoe@company.com;dsmith@company.com • To use a Scheduler job variable to define a notification recipient, enter the variable name with curly brackets. You can combine regular email addresses and variables, such as: {JobOwner.EmailAddress}; jdoe@company.com <p>By default, the notification is configured to be sent to the user who executed the job, using the variable {CurrentUser.EmailAddress}.</p> <p>The entries in the To field must be valid email addresses, or Scheduler job variables that will resolve to valid email addresses. Currently, it is not supported to list user or role names, or to look up email addresses from Axiom Security.</p> <p>NOTE: When using Send email notification to different email addresses when the job has errors or succeeds, this user will be notified if the job completes successfully (including partial success), but not if the job fails. Job failure notifications are sent to the To (on error) recipients.</p> |

| Item | Description |
|--------------|---|
| From | <p>The email address that the message is sent from. This can be something like <code>axiomscheduler@company.com</code>, so that the recipient can easily tell that the message has been generated by Scheduler.</p> <p>By default, this is set to the Scheduler "from" email address as defined in the system configuration settings, using the system variable <code>{Scheduler.FromEmailAddress}</code>.</p> <p>NOTE: For installations that are using subsystems, the system variable <code>{Scheduler.FromEmailAddress}</code> may resolve to a subsystem administrator email address instead of the Scheduler "from" email address.</p> |
| Subject | The subject of the message. By default, this is set to "Axiom Scheduler Notification." |
| User Message | Optional body text for the email. This text is included in addition to the Scheduler auto-generated text regarding the job status. |

If **Send email notification to different email addresses when the job has errors or succeeds** is enabled, the following additional options are available:

| Item | Description |
|--------------------|--|
| To (on error) | <p>The email address(es) to receive the notification email when the job result is Failed. Separate multiple addresses with a semicolon.</p> <p>This user only receives a notification if the job fails. If the job result is Success or Partial Success, this user will not receive a notification (only the To user will).</p> |
| Subject (on error) | The subject of the job failure message. By default, this is set to "Axiom Scheduler Notification." |

Job variables can be used in all notification settings.

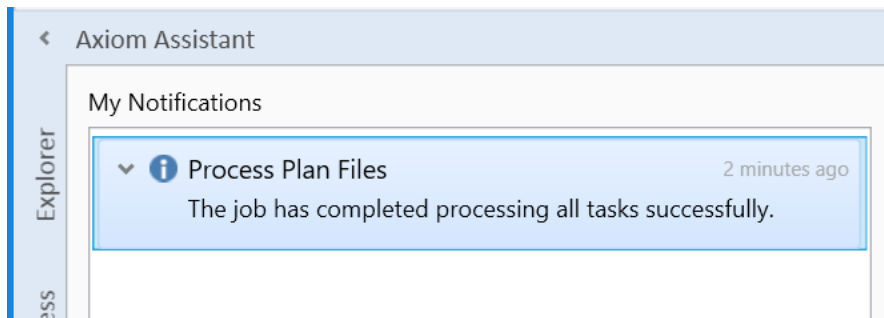
► Application notifications for Scheduler jobs that are run manually

If you run a Scheduler job manually, you can receive a notification within the application to let you know the status of the job. This notification will display in the Notifications task pane of the Desktop Client, and in the Notifications panel of the Web Client. This notification works as follows:

- The in-application notification is only sent if the Scheduler job is run manually using the **Run Now** option in Scheduler (or by using an equivalent "run now" action within a product-specific web page). In-application notifications are *not* sent if the job is run via a scheduling rule or an event handler.

- The in-application notification honors the Notification settings defined for the job to determine whether the notification is sent. For example, if the job is set to **None**, then the in-application notification is not sent. If the job is set to **Send all email notifications**, then both an email notification and an in-application notification will be sent when the job completes.
- The in-application notification only reports the status of the job—success, failure, or partial success. It does not contain any error or success details, and does not include any messaging as defined in the Notification settings for the job. For more information, view the job results within the Scheduler dialog in the Desktop Client, or the Scheduler page of the Web Client.
- The in-application notification is always sent to the user who ran the job manually.

NOTE: If the job is configured to **Send email notification to different email addresses when the job has errors or succeeds**, this is treated as **Send all email notifications** for purposes of sending the in-application notification. The user who ran the job will be notified when the job is completed, regardless of the job status.



Example success notification

Creating event handlers for a job

You can create user-defined event handlers in a job, for the purposes of automatically triggering the job for execution when the event name is called by another feature. Axiom supports several features that can be used to raise an event:

- The RunEvent function and command
- File Group triggers
- The Raise Event Scheduler task

Event handlers are defined by name. Multiple jobs can have an event handler with the same name. When that event handler is called, it will affect all jobs that contain the event handler with the matching name.

To create an event handler in a job:

1. In the **Scheduler** dialog, open a job to edit or create a new job.

2. In the left-hand side of the job, select **Event Handlers**.
3. On the **Job** tab of the ribbon, in the **Event Handlers** group, click **Add**.

A new event handler is added to the job.

4. Double-click the **Event Name** field so that the field becomes editable, and then type the desired event handler name.

For example, if the event handler will be used to trigger Process Plan Files jobs, you might name the event handler ProcessPlanFiles.

This event name is the name that will be used in features such as RunEvent to trigger this job for execution.

5. In the **Execute As** field, select one of the following to determine the user identity that will be used to run the job when it is executed via the event handler:

- **Owner:** The job will be run under the identity of the job owner.

The job owner is the user who last saved the job. If you are not sure who the current job owner is, you can check the **Job Variables** tab. The current job owner is listed in the **System defined values** section.

- **Requester:** The job will be run under the identity of the user who triggered the event handler.

By default, event handlers are set to run as the **Owner**. You should carefully consider this option as it may affect whether the job can be run and how the job is run.

For example, if the event handler is set to **Requester**, but the user who triggered the job does not have access to the file group specified for a Process Plan Files task, then the task will fail.

This may be the desired outcome—you may want the job to be dependent on the user's rights, and therefore you should specify **Requester**. On the other hand, you may want the job to run in the same way every time, regardless of the user that triggers the job. In that case you should specify **Owner**.

By default, the event handler is set to **Active**, which means it will be found by any process that triggers the event handler. If you want to temporarily exclude this job from event handler processing, you can clear the Active check box.

► Associating an event handler with a file group

If a Scheduler job with an active event handler is stored in a file group Utilities folder, then the event handler is associated with that file group. When using RunEvent to trigger jobs for execution, you can optionally specify a file group context so that only event handlers associated with that file group (or no file group) are considered.

In order to store a Scheduler job within a file group, you must first create and save the job within the Scheduler Jobs Library. Then, you can use Axiom Explorer to move the job from the Scheduler Jobs Library to the file group Utilities folder. If the file group and its utilities are later cloned using any process—such as regular file group cloning, scenario creation, or file group rollover—then the event handler is also cloned and will be associated with the new file group.

► Deleting an event handler

If you no longer need an event handler, you can delete it from the job. Select the event handler and then click **Remove Selected**. You can also **Clear All** to remove all event handlers from the job.

User-defined event handlers display along with the system event handlers in the **Event Handlers** tab (**Service > Event Handlers**). If you right-click a user-defined event handler in this location and select **Remove event handler**, it does not delete the event handler from any jobs that use it, but it does set the event handler to inactive.

Using job variables

You can use job variables within a Scheduler job, to define the value of the variable when the job is run. Job variables are managed in the **Job Variables** section of the job.

There are two types of variables:

- **User-defined variables:** You can create a variable and then use it within any job or task setting that supports variables, but only within that particular job. The primary use for user-defined variables is to run a job via RunEvent (either the function or the command), and pass in a variable value at that time.
- **System variables:** Axiom provides a number of system variables that can be used within relevant job and task settings. For example, instead of specifying a "hard-coded" email address for the job notification, you can use a system variable to specify that whoever ran the job should receive the notification.

In all cases, to use a variable within a job or task, enter the variable name into the desired setting, enclosed in curly brackets. For example: {variable}

TIP: If you want to use a variable in a job, you can right-click the variable and then select **Copy variable name to clipboard**. Navigate to the setting where you want to use the variable, and then paste it into the setting (the curly brackets are added automatically).

At the job level, variables can be used in any of the **Notification Message Content** settings in the **Notification** tab. At the task level, in general, variables can be used in any task field that accepts typed user input.

When the job is run, the variable values used for the job display in the job results under **Job Values**, and also in the email notification (if applicable).

► User-defined variables

User-defined variables are created in the **Job Variables** tab. You define the name of the variable (without brackets), and if desired, define a default value for the variable.

When the job is run, the user-defined variable will be replaced with a value as follows:

- If the job was scheduled using RunEvent (function or command), and RunEvent sent a name / value pair that matches the name of the job variable, that value is used.
- If the job was scheduled as a result of a file group trigger, and the trigger has a defined variable that matches the name of the job variable, that value is used. Multiple values are returned as a comma-separated list.
- If the job contains a Process Document List task or a Process Plan Files task with a defined post-processing variable that matches the name of a job variable, that value is used after that task has been processed.
- Otherwise, the default value defined in the Job Variables tab is used.

If the value is blank, then the job or task setting using the variable will be evaluated as blank. If the setting cannot be blank, then an error will result when the job is executed.

► System variables

The available system variables are listed at the bottom of the **Job Variables** tab. Most of these variables relate to user names and addresses, for use within the job notification settings.

When the job is run, the system variable is replaced with the applicable system value.

The following values are available:

| Variable | Description |
|---|--|
| <code>{CurrentUser.EmailAddress}</code> | Returns the current user's email address, login name, or full name. |
| <code>{CurrentUser.LoginName}</code> | The current user is the user identity under which the job is currently being run. Generally, this is the user who executed the job. If the job was executed via an event handler and the event handler is set to owner, then the current user will be the job owner. |
| <code>{CurrentUser.FullName}</code> | |
| <code>{JobOwner.EmailAddress}</code> | Returns the job owner's email address, login name, or full name. |
| <code>{JobOwner.LoginName}</code> | The job owner is the user who last saved the job. |
| <code>{JobOwner.FullName}</code> | |

| Variable | Description |
|---|--|
| <code>{Scheduler.ConfiguredFromEmailAddress}</code> | Returns the system's default "from" address, as defined in the system configuration settings. |
| <code>{Scheduler.FromEmailAddress}</code> | <p>This returns a value as follows:</p> <ul style="list-style-type: none"> • If the current user belongs to a subsystem, this returns the subsystem administrator's email address. • If the current user does not belong to a subsystem, this returns the default configured "from" address. |
| <code>{CurrentSubsystem.AdminEmailAddress}</code> | <p>Returns the email address of the subsystem administrator for the subsystem that the current user belongs to.</p> <ul style="list-style-type: none"> • If the subsystem has multiple administrators, the email is sent to the first administrator. • If the user belongs to multiple subsystems, the first returned subsystem for the user will be used. No specific logic is applied to determine the "correct" subsystem for any particular job. • If the user does not belong to a subsystem, then no email address is returned. |
| <code>{EventHandler.EventName}</code> | Returns the name of the event handler that caused the job to be scheduled, if applicable. Otherwise the variable returns blank. |

| Variable | Description |
|--|--|
| <code>{NotificationAddress}</code> | <p>Returns the notification address defined for the plan codes that triggered a Scheduler job.</p> <p>This variable only applies when the job is executed as a result of a file group trigger, and only if the optional notification address settings are defined for the file group. Otherwise, no email address is returned.</p> |
| <code>{Task.CurrentIterationValue}</code> <code>{Task.IterationNumber}</code> | <p>Returns the current iteration value and the current iteration number. These variables only apply when using the Iteration feature for a task.</p> <p>For more information, see Using iterative task processing.</p> |

Processing tasks in parallel

Each Scheduler job can have multiple tasks. By default, each task in the job is processed sequentially, in the order that the tasks are listed in the job.


If desired, you can configure tasks so that they are processed concurrently (in parallel) instead of sequentially. If appropriate, this may speed up the processing of the job.

► Configuring tasks for parallel processing

In order to process tasks in parallel, the tasks must be configured to run as subordinate jobs (sub-jobs). To do this, edit the following settings in the **Task Control** section for each task:

- Select **Create a Subordinate Job for this Task**.
- Ensure that the following setting is *not* selected: **Wait for all Subordinate Jobs to complete before proceeding to the next Task**.

In the following example, if all four tasks are configured to be run as subordinate jobs, then they can be run in parallel (depending on the available Scheduler threads).



Scheduler task configured to run as a subordinate job to enable parallel processing

► How parallel processing works

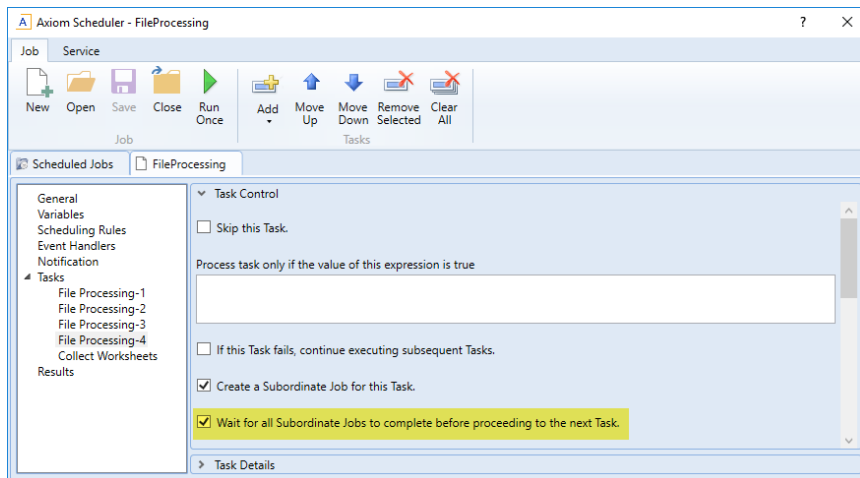
When a task is configured to execute as a subordinate job, then it is not processed within the "parent" job. Instead, a sub-job is created for the task. The sub-job joins the Scheduler queue and is eligible for processing according to the normal Scheduler processing rules. For more information, see [Processing priority for scheduled jobs](#).

For example, imagine that you have a job with four tasks, and these tasks are not dependent on each other. If you use the default settings, Scheduler takes the first task in the list and starts processing. The second task is not started until the first task is complete, and so on.

If instead you configure each task as a sub-job, then when the "parent" job is processed, it will create four sub-jobs. If two Scheduler threads are available for processing, then two of the sub-jobs are processed at the same time. If four Scheduler threads are available, then all four sub-jobs are processed at the same time. Once all of the sub-jobs are complete, the parent job is completed, and its status reflects the overall status of all of the sub-jobs.

If tasks are dependent on each other, then you should not process them as sub-jobs, or you should use the **Wait** setting as appropriate. For example, imagine that the first four tasks in the job can be run in any order, but the fifth task must be processed last. In that case, you can configure the first four tasks to run as sub-jobs, but on the fourth task you must enable **Wait for all Subordinate Jobs to complete before proceeding to the next Task**. This will cause Scheduler to wait for all sub-jobs to finish before it proceeds to the fifth, final task.

In the following example, the file processing tasks are configured as sub-jobs so that they can be run in parallel. The last file processing task is configured to wait, so that all of the file processing tasks will be finished before the file collection task begins.



Scheduler task configured to wait for all subordinate jobs to complete

► Using an Echo task to control jobs with sub-jobs

By default, if a sub-job fails, this failure status will not prevent the parent job from continuing to process tasks. If you want the parent job to stop processing on sub-job failure, you can use an Echo task to do this.

If the **Task Control** option **Wait for all Subordinate jobs to complete before proceeding to the next task** is enabled for an Echo task, then the Echo task will inherit the status of the sub-jobs when they are completed. If one of the sub-jobs fails, then the Echo task will inherit that failed status. By default, this means that the job will not process any further tasks due to the sub-job failure.

In the previous example, if one of the file processing tasks being run as a sub-job fails, this will not stop the final file collection task from being processed. To work around this, you can place an Echo task after the file processing tasks and configure it to wait, and then if any of the file processing sub-jobs fail the Echo task will fail. This task failure within the parent job will stop the final file collection task from being processed. (When using this configuration, the fourth file processing task does not need to be configured to wait, because the job will wait on the Echo task instead.)

Using iterative task processing

You can configure a Scheduler task to use iterative processing, so that the task is repeated multiple times using a designated list of values. Each iteration of the task uses a different item in the list, until all items have been processed.

When you enable iterative processing for a task, you define the list of values by specifying a table column and an optional filter. The task will then be processed for each unique item in the table column. You can reference the column values within the task properties by using a built-in Scheduler job variable. As each iteration of the task is processed, the variable is replaced with the column value for the current iteration. Using this approach, the task can dynamically change for each iteration.

For example, you may have an import that you want to perform for four different entities in your organization. The import configuration is exactly the same except that the source file or query is different for each entity. If the import uses entity as a variable, then you can set up a single import task and configure it to iterate over the list of entities. Each iteration uses a different entity name or code, which you can pass into the import variables so that the import uses the correct source file or query for the current entity.

► Enabling iterative processing

Iterative processing is enabled in the Task Control properties of the task. Select the task within the Scheduler job, then click **Task Control** to expand that section. Any task can use iterative processing, though it is more useful for certain task types such as Import ETL Package.

Complete the following properties in the **Iteration** section of the Task Control properties.

| Item | Description |
|---|---|
| Iterate this Task | Specifies whether iterative processing is enabled for the task. If enabled, then the task will be performed N times, where N is the number of unique items in the specified iteration column. Job variables can be used to apply the current iteration value and iteration number to the task. |
| Create a Subordinate Job for each iteration | <p>Specifies whether each iteration is processed as a separate subordinate job. By default, this is disabled, which means that all iterations are processed sequentially within the overall subordinate job created to process the iterations.</p> <p>If enabled, then each iteration is processed as a separate subordinate job, enabling concurrent execution of multiple iterations. This option should only be enabled if the order of iteration processing is not important.</p> |
| Column | <p>The column that contains the values to iterate over. Use Table.Column syntax to specify the column. Multiple-level lookups can be used.</p> <p>For example, if you specify <code>Dept.Region</code>, then the task will be processed once for each unique region value in the column (after applying any filter to limit the list of values).</p> |
| Group By | <p>Optional. By default, the group by column is the same as the iteration column, so that the task is processed once for each unique value in the iteration column. However, if needed, you can specify a different grouping level.</p> <p>You can use any column or columns that would be valid as the "sum by" level for an Axiom query, where the primary table is the table specified for the iteration column.</p> |

| Item | Description |
|----------|--|
| Order By | <p>Optional. By default, the values are sorted based on the iteration column, in ascending order. You can specify a different sort column, or use the same sort column but change the order to descending.</p> <p>The sort order is ascending unless the keyword <code>desc</code> is used to specify descending order. For example:</p> <pre>Dept.Dept desc</pre> |
| Filter | <p>Optional. A filter criteria statement to limit the list of values for the iterative processing. You can use any filter that is valid against the source table (the table of the iteration column).</p> |

When iterative processing is enabled for a task, the iterations are always processed within a subordinate job. Therefore, enabling the Task Control option of **Create a Subordinate Job for this Task** is unnecessary.


If your job has multiple tasks, and you want the tasks after the iterative task to wait for all iterations to complete before executing, then you must enable the following Task Control option for the iterative task: **Wait for all Subordinate Jobs to complete before proceeding to the next Task**.

► Configuring the task to change for each iteration

In order for the Scheduler task to apply the current iteration value to each iteration, you must use the built-in iteration variables within the task. These variables are job variables, and can be used like any other job variable. The following variables are available:


| Variable | Description |
|---|--|
| <code>{Task.CurrentIterationValue}</code> | Returns the current value from the iteration list. |
| <code>{Task.IterationNumber}</code> | Returns the number of the current iteration. |

To continue the previous example, imagine that you are setting up an import for iterative processing by entity. To define the list of entities, you set up the Iteration settings in the Task Control section like the following:



This example will iterate over the list of entities in the Entities column, limited to only those entities where the Costing column is set to True. If this resolves to 4 entities, then the task will be processed 4 times, once for each entity.

The import is configured with a variable `{Entity}`, which it uses to process the correct entity source file. In order to pass the current task iteration value to the import variable, you can use the job variable `{Task.CurrentIterationValue}` in the import task settings. For example:




When the first iteration is performed, the `{Task.CurrentIterationValue}` will be resolved as `Entity_1`, so the import will be processed using `Entity_1` as the value for the `{Entity}` import variable. For the second iteration, the value `Entity_2` will be used, and so on. Using this approach, the import will be processed for all entities in the iteration column.

Conditionally processing tasks in a job

You can configure a task so that it is only processed if a particular condition is met. This feature is configured in the task settings, in the **Task Control** section, under **Process task only if the value of this expression is true**.

To enable conditional processing, you must specify a logical expression that will resolve to either true or false when the job is executed. If true, then the task is processed as normal. If false, then the task is skipped.



The logical expression is evaluated using an IF function on the Scheduler server as follows:

```
=IF(Expression, 1, 0)
```

You can enter any expression that would be valid in an IF function. You can use Excel functions, Axiom functions, and Scheduler job variables in the expression. If you use a job variable, it must be placed in quotation marks unless you expect the variable value to be resolved and evaluated as a number.

If the task is not processed because the condition resolves to false, this is not considered a failed task. If there are other tasks in the job, they will be processed. If you want an entire job to be conditional, you can do either of the following:

- Repeat the condition in each individual task settings. Keep in mind that the condition will be evaluated for each individual task, which means that if it is possible for the condition to change in between tasks, some tasks might be processed while others aren't.
- Use the condition on a Raise Event task that then triggers another job for processing. For more information, see [Raise Event task](#).

► Examples

The following are some example expressions for conditional processing:

```
GetData("Process", "FG='2012 Budget'", "SystemStatus")=1
```

If this GetData function returns 1, the expression resolves to true and the task is processed. If not, it is false and the task is skipped.

```
AND (" {EventHandler.EventName}"="ProcessPlanFiles", {Dept}=1000)
```

If this job was triggered for execution by the ProcessPlanFiles event handler, and if the job variable Dept resolves to 1000, then this expression is true and the task is processed. Note that in the first part of the expression, the event handler variable will return a string value so it must be placed in double quotation marks. In the second part of the expression, the department variable will return a number so it is not placed in quotation marks.

```
AND (Day (Now () ) <=7, Weekday (Now () ) =2)
```

This expression will return true if it is the first Monday of the month, otherwise it will return false.

Run another Scheduler job from within a Scheduler job

Scheduler jobs have two ways to run another Scheduler job:

- **Raise Event task:** This task uses an event handler name to trigger one or more Scheduler jobs for execution. The jobs triggered by the event handler are run independently from the job containing the Raise Event task.
- **Run Scheduler Job task:** This task runs a specified Scheduler job as a subordinate job within the current "parent" job. Essentially, the tasks in the target job are run within the parent job, which means that other tasks in the parent job can reference the results of those tasks.

The decision of which task to use depends on several factors, but the most important is whether the Scheduler jobs are independent or dependent. If the first job is dependent on the execution of the second job, then you must use the Run Scheduler Job task. When the Raise Event task is used, the triggered jobs are run independently.

The following chart details some comparison points between the two tasks:

| Comparison | Raise Event | Run Scheduler Job |
|---|-------------|-------------------|
| Can pass variables to target job | Yes | Yes |
| Can execute target job as requester or owner | Yes | No |
| Can wait for target job to complete before continuing | No | Yes |
| Can use results of target job in subsequent tasks | No | Yes |

► Using the Raise Event task

The Scheduler Raise Event task is typically used when you need to trigger another job for execution once the current job is complete. The Scheduler job(s) triggered by the event handler are added to the schedule and then executed independently from the current job. For example, you might place the Raise Event task at the end of the task list, so that all tasks in the current job must complete successfully before the Raise Event task is run.

Because the jobs triggered by the event handler are run independently, you cannot perform additional tasks in the current job that depend on the results of the triggered jobs. The current job will not wait for the triggered jobs to be run.

Event handlers can be configured to run a job as either the job owner or as the requester (meaning the user that triggered the event). This allows some additional flexibility in how the triggered jobs are run. For example, you may want to trigger a job that needs to be run using administrator permissions. As long as the event handler is configured to run as job owner (and the owner is an administrator), then that triggered job will always have the necessary permissions, regardless of the user who is running the job with the Raise Event task.

► Using the Run Scheduler Job task

The Run Scheduler Job task is typically used when you need to run another job and then perform additional tasks once that job is complete. Because the target job is run as a subordinate job within the "parent" job, the parent job can wait for the "child" job to complete before it continues processing tasks. Later tasks in the parent job can reference the results of the completed child job, such as querying data saved from the child job, processing plan files created by the child job, and so on.

Keep in mind that if you want the parent job to stop processing tasks if the target job fails, then **Wait for all Subordinate Jobs to complete before proceeding to the next Task** must be enabled for the task and **If this Task fails, continue executing subsequent Tasks** must be disabled for the task. This is the default configuration.

The child job is run using the same user permissions as the parent job. The user running the parent job must have the appropriate permissions to complete all tasks in both the parent job and the child job.

► Chaining multiple Scheduler jobs

You can "chain" multiple Scheduler jobs together using either approach. For example, you may have three Scheduler jobs that you want to run, in a particular order.

You can place multiple Run Scheduler Job tasks in a parent job, where each task triggers a separate job. Because these tasks run as subordinate jobs, and the parent job can wait for each child job to complete, it is easy to run the jobs in order. However, if you want to stop processing the jobs if one fails, there is no built-in way to do that (as discussed in the previous section). You would need to set up the Run Scheduler Job tasks to run or not based on a condition, where the condition tests some result from the previous job.

To chain jobs using Raise Event, the last task in each job can be a Raise Event task. Each job will perform its tasks and then trigger the next job in the chain. When using this approach, the chain automatically stops if failure occurs, because if a task in the job fails then the job stops and will not proceed to the Raise Event task. The disadvantage of this approach is that the jobs cannot also be run separately, unless you manually disable the Raise Event tasks or configure the Raise Event tasks to not run based on a condition.

Running a job

If a job is saved with an active scheduling rule, then the job is automatically placed on the schedule to be run according to that rule. Each time the job is run according to the rule, it is run as the current job owner (unless it is a system job, in which case it is run as the System identity).

However, you can also choose to run a job manually. If you run a job manually, the job is added to the **Scheduled Jobs** list with a start time of now, to be processed according to its [job priority settings](#). The job will be run using your user identity (again, unless it is a system job).

Running a job manually does not impact any scheduled executions of the job as determined by scheduling rules. For example, if a job is scheduled to be run at 10:00 PM tonight, and you run the job manually at 2:00 PM, the job will still be run as scheduled at 10:00 PM.

To manually run a job:

1. In the [Scheduler dialog](#), in the **Job** tab, click **Open**.

The **Axiom Explorer** dialog opens, showing the Scheduler Jobs Library only.

2. Select the job and then click **Open**.

The job opens in the **Scheduler** dialog. Make sure the job is the active tab in the navigation pane (the most recently opened tab is the active tab by default).

3. In the **Job** tab of the ribbon, click **Run Once**.

A confirmation message informs you that the job has been placed on the schedule.

NOTE: When you click **Run Once**, any unsaved changes to the job are automatically saved. This save will designate you as the job owner (if you are not already the job owner).

You can also run jobs manually using the Scheduler area of the Web Client. For more information, see [Running a job manually in the Web Client](#).

Using RunEvent to execute a Scheduler job

Using RunEvent, you can trigger the execution of a Scheduler job from various contexts, such as within Axiom files, task panes, or Axiom forms. There are two different versions of RunEvent:

- **RunEvent function:** The RunEvent function can be used in Axiom files to trigger the execution of a Scheduler job from a spreadsheet.
- **RunEvent command:** The RunEvent command can be used in task panes or Axiom forms to trigger the execution of a Scheduler job.

Both the function and the command work in the same way and use similar parameters. Some limitations apply depending on the context where RunEvent is being used. It is assumed that an administrator (or a power user with the necessary rights) sets up the desired jobs within Scheduler, and then sets up RunEvent in the appropriate context so that end users can trigger it.


The end user who triggers the job using RunEvent does *not* need to have file permission to the job or any access to Scheduler. The job itself can be configured to execute its tasks using the permissions of the job owner or using the permissions of the end user who triggered the job (the requester). If the job is run as the requester, then the end user must have the appropriate permissions to the files impacted by the job (for example to the target file for File Processing, or to the target file group and plan files for Process Plan Files).

NOTE: You can also use the [Raise Event](#) Scheduler task to trigger the execution of a Scheduler job from a different Scheduler job. This works in a similar manner as the RunEvent features.

► Setting up a Scheduler job for RunEvent

All uses of RunEvent require the same job setup in Scheduler:

- The job that you want to execute via RunEvent must already be created within Scheduler. When setting up the job, consider items such as the notification settings. Do you want the notifications to go to the user that executed the job, or to the job owner, or both?
- The job must contain an event handler that will be used to trigger the job execution. When creating the event handler, consider whether you want the job to run as the job owner, or as the requester (the user who clicks on the RunEvent function). This may impact email notifications and determines the user rights under which the job will run.



For more information, see [Creating event handlers for a job](#).

- Optionally, the job can use variables. Variable values can be defined in the RunEvent function or command and passed to the job. You would do this if aspects of the job need to be dynamic; for example, if you want to use a different filter depending on which user is running the job or based on a user selection in the file. For more information see the *Variable example* section below.

► Setting up RunEvent

RunEvent uses the following properties to trigger Scheduler jobs:

- The event handler name that identifies the Scheduler job(s) to trigger for execution.

- An optional confirmation message to present to the user before proceeding with the event. Not available when using the command within an Axiom form.
- An optional success message to present to the user after the event has been raised.
- An optional file group context to target the job execution to only event handlers that are associated with a particular file group (or no file group). For the RunEvent command, this is an optional parameter. For the RunEvent function, the current file group context is automatically applied if the file with the function belongs to a file group.
- If variables are being used, one or more variable names and values to pass to the Scheduler job. This is available in all contexts, however, task panes do not currently support the ability to determine the variable values dynamically.

The following shows an example RunEvent function for use in an Axiom file:

```
=RunEvent("Click here to process plan files","ProcessPlanFiles","You are about to process plan files for the "&C3&" region. Do you want to continue?",,"filter = "&C4)
```

The first parameter defines the display text for the function, while the second parameter specifies the event handler name. In this example we have also defined a custom confirmation message for the user and a variable value to pass a filter to the job. The following screenshot shows the function in the spreadsheet:

| | B | C | D | E | F |
|---|---|--|---|---|---|
| 1 | | | | | |
| 2 | | | | | |
| 3 | | North | | | |
| 4 | | Dept.Region='North' | | | |
| 5 | | | | | |
| 6 | | | | | |
| 7 | | Click here to process plan files | | | |
| 8 | | | | | |

The next screenshot shows a RunEvent command set up on a Button component for an Axiom form. You can see the same event name and the filter variable also being read from a sheet location.

Shortcut Properties

Choose Shortcut

Shortcut Target
 command://RunEvent ...

Shortcut Parameters

Event Name: ProcessPlanFiles

Result Message: Process Plan Files has been scheduled.

File Group:

Processing Step: Form - Before Processing

Event Variables

| Filter | |
|--------|--------------|
| | [Sheet1!A15] |


OK Cancel

► RunEvent behavior and user experience

The behavior and user experience for RunEvent depends on the context and whether you are using optional custom messages.

- The user starts the process by doing one of the following:
 - Double-clicking the RunEvent function in the spreadsheet.
 - Clicking the Button component that is configured for RunEvent in the Axiom form.
 - Double-clicking the RunEvent item in the task pane.

- A confirmation prompt displays to the user, asking them to confirm that they want to proceed. The user can click **Yes** to proceed, or **No** to cancel. Default text is used if no custom text is defined in the RunEvent properties.



NOTE: This step does not apply when executing RunEvent from an Axiom form. The Axiom form context does not support a confirmation message. However, you can configure the Button component to display a confirmation message before executing the RunEvent command.

- Axiom checks the list of event handlers in Scheduler to see if any match the specified event handler name in RunEvent. This check works as follows:
 - If RunEvent has a file group context, then Axiom only tries to match with event handlers that are associated with the same file group, or with no file group. Any event handlers associated with a different file group are ignored. The RunEvent command has a file group context if a file group is specified in the shortcut parameters, whereas the RunEvent function automatically has a file group context if the file with the function belongs to a file group.
 - If RunEvent does not have a file group context, then all event handlers are eligible to match.

If any matching event handler names are found, then all Scheduler jobs that reference the event handler are triggered for execution. If multiple jobs reference the matching event handler, then all of those jobs will be executed.

If variable values are defined in the RunEvent properties, those values are passed to the job and are used when the job is executed.

- A confirmation message displays to the user as follows:
 - If no jobs were found that contained the specified event handler, the user is notified that no jobs were found.
 - If jobs were placed on the schedule, the user is notified that the specified event was scheduled. Default text is used if no custom text is defined in the RunEvent properties.

NOTE: If executing RunEvent from an Axiom form, this message displays in the bottom left corner of the form, not in a separate message dialog.


When the job is finished processing, email notifications are sent according to the settings in the job.

► Variable example

When using RunEvent to execute a Scheduler job, you can pass a variable value to the job. For example, imagine that you want to execute a Process Plan Files job, and you want to send a filter value to the job.


Step 1: Set up the variable in the job

The first step would be to create a job variable in the job, and then use the job variable in the filter setting.



Example job variable

NOTE: In this case, we have defined a default value for the filter variable (1=0) that does not result in any plan files. This is because we do not want to process any plan files unless a filter is provided by RunEvent. If we left the default value blank, that would mean all plan files would be processed if no filter was provided by RunEvent.



Example use of variable in job settings

For more information on Scheduler job variables, see [Using job variables](#).

Step 2: Configure RunEvent to use the Variable

Now that the job is set up to use the filter variable, you must configure RunEvent to pass in a value for that variable. If you are using the RunEvent function in an Axiom file, you use the following syntax within the function parameters:

```
variablename=variablevalue
```

These name / value pairs can be placed in the RunEvent function starting in the fifth parameter of the function. If you have two name / value pairs to pass to the job, you can use the fifth and sixth parameters, and so on.

For example, to pass the filter `DEPT.Region='North'` to the job, the RunEvent function would be constructed as follows:

```
=RunEvent("Double-click to process plan  
files", "ProcessPlanFiles", , , "filter=dept.region='North'")
```

More likely, you would be reading the variable value from another place in the sheet, so the function would look something like:

```
=RunEvent("Double-click to process plan  
files", "ProcessPlanFiles", , , "filter="&C4)
```

Where the filter value is read from cell C5.

When the job is executed by use of this RunEvent function, the value `DEPT.Region='North'` will be placed in the **Plan File Filter** box of the Process Plan Files task, and the job will be run using that filter.

When using RunEvent in an Axiom form, the variables and their values are defined in the Event Variables section. In this context you place the name of the variable in the left-hand box, and then in the right-hand box you enter the cell reference (in brackets) where the variable value will be read.



The Event Variables section is also present when configuring RunEvent for use in a task pane, however in this context the only option is to "hard-code" the values in the RunEvent properties.

Job properties

This topic is a reference for the settings that can be defined for a Scheduler job.

► General

This section defines general settings for the job.

| Item | Description |
|----------------------|---|
| Description | <p>Optional. The description of the job.</p> <p>The job description can also be edited in Axiom Explorer, in the Scheduler Jobs Library.</p> |
| Job Restart Behavior | <p>Specifies whether and how the job should be restarted if it is interrupted prior to completion. Select one of the following:</p> <ul style="list-style-type: none">• Do not reschedule this job. In this case, you must manually reschedule the job if it needs to be run before its next scheduled execution.• Restart the job from the first task. The entire job is run again, even if some of the tasks were completed successfully before the job was interrupted.• Resume the job beginning with the first uncompleted task. (Default) The job resumes and only the uncompleted tasks are run. <p>A job would be interrupted if the Scheduler server processing it was restarted, or if the Scheduler service on the server was stopped or restarted, or if the Scheduler server was disabled from the Servers tab (Service > Servers) of the Scheduler dialog.</p> |
| Job Results Cleanup | <p>Specifies whether historical job results are purged when the job is run.</p> <p>To purge job results:</p> <ol style="list-style-type: none">1. Select Purge historical job results whenever this job runs.2. In Number of days to keep results for this job, specify the number of days to keep when purging results. By default this is set to 0, which means all job results will be purged except the result for the current job execution. <p>A day is counted as 24 hours from the time the cleanup task is executed. So if you specify 1 day, and the task is run at 11:00 PM on Tuesday, then all results prior to 11:00 PM Monday are purged.</p> <p>If this option is not selected, then historical job results remain in the database until the system's Purge System Data task is run.</p> |

| Item | Description |
|--------------------|--|
| Priority Elevation | <p>Specifies the priority of the job in the scheduled jobs queue, within the job's priority category. Select one of the following:</p> <ul style="list-style-type: none"> • Default: (Default) This job is run on a "first come, first served" basis. The total number of jobs that can be run at one time is determined by the configured number of Scheduler threads for the installation. • Reduced: The job is designated as a low priority job, and remains at the bottom of the queue until other jobs with Default and Elevated priority have been run. • Elevated: The job is designated as a high priority job, and is moved to the top of the queue to be run before Default and Reduced priority jobs. • Interrupt: The job is run immediately, regardless of any jobs currently waiting in the scheduled jobs queue, and regardless of whether any Scheduler threads are currently available to run it. If no Scheduler threads are currently available, a new thread is created, even if this exceeds the configured thread limit for the installation. <p>Job execution order also depends on the priority category of a specific job execution. See Processing priority for scheduled jobs.</p> |
| Mark as System Job | <p>Specifies whether the job is run as a system job. Only administrators can edit this check box.</p> <p>If this check box is selected, the job is run under the "Scheduler Service" system identity instead of a user identity, and the job is run by the system Scheduler server which operates on the Axiom Application Server.</p> <p>Generally, this check box should only be selected for system "support" tasks that should not depend on individual user rights. This check box is not available if the job contains non-system tasks (generally, spreadsheet-related tasks).</p> <p>For more information, see System jobs.</p> |

| Item | Description |
|---|---|
| Put the system in 'admin only' mode during this job | <p>If this option is selected, then the system will be placed into administrator-only mode at the start of the job, and then placed back into full access mode when all tasks are completed (including any sub-jobs). This is the same behavior as going to Manage > Security > System Access and selecting Administrators Only.</p> <p>NOTES:</p> <ul style="list-style-type: none"> You should make sure that any jobs using admin-only mode do not overlap. For example, imagine that job A starts and places the system in admin-only mode. While job A is still running, job B starts and finishes. If job B also uses admin-only mode, then when job B finishes the system will be placed back into full access mode, meaning the remainder of job A will be processed in full access mode. Any job using admin-only mode must be run by an administrator. <p>Generally speaking, any job set to run using admin-only mode should be run at a time when no end users will be logged into the system and no other Scheduler jobs will be running.</p> |

► Job Variables

This tab has two sections for job variables:

- In the **Job values** section at the top of the tab, you can manage user-defined variables for use in the current job.

To add or remove variables, use the **Add**, **Remove Selected**, or **Clear All** commands in the **Job Variables** group of the **Job** tab. This group is only available when you have selected the **Job Variables** section in the left-hand side of the job.

When creating user-defined variables, do not add curly brackets to the variable name. Curly brackets are only required when you use the variable in a job or task setting.

- In the **System defined values** section at the bottom of the tab, you can view the system variables available for use in the job.

You can right-click any variable in this section (user-defined or system-defined) and select **Copy variable name to clipboard**. You can then navigate to the setting where you want to use the variable, and then paste it. The variable will be pasted with the necessary curly brackets.

For more information, see [Using job variables](#).

► Scheduling Rules

Each row in this section defines a scheduling rule for the job. Jobs will be automatically scheduled according to the settings in this section.

To add or remove scheduling rules, use the **Add**, **Remove Selected**, or **Clear All** commands in the **Scheduling Rules** group of the **Job** tab. This group is only available when you have selected the **Scheduling Rules** section in the left-hand side of the job.

For more information, see [Defining scheduling rules for a job](#).

| Item | Description |
|-------------|--|
| Active | Specifies whether the scheduling rule is active. If this check box is not selected, then the rule is ignored for purposes of scheduling the job. |
| Starting On | <p>Optional. Specifies the earliest date and time for the scheduling rule to take effect.</p> <p>If you want the job to run one time only, set Starting On and Ending On dates to the same date/time.</p> |
| Ending On | Optional. Specifies the expiration date and time for the scheduling rule. Once this date is past, no further executions will be scheduled for this rule. |
| Day of Week | <p>Specifies the day(s) of the week that you want the job to be run:</p> <ul style="list-style-type: none">• * (Default): The job will be run on all days within the start / end range.• 0-6: The job will be run on the specified day or days, where 0 is Sunday and 6 is Saturday. Use a comma or a hyphen to separate multiple days (hyphen if the days are contiguous, commas if not). |
| Hours | <p>Specifies the time of day (hours) that you want the job to be run, in relation to the specified days:</p> <ul style="list-style-type: none">• * (Default): The job will be run on all hours.• 0-23: The job will be run on the specified hour or hours, where 0 is midnight and 23 is 11:00 PM. Use a comma or a hyphen to separate multiple hours (hyphen if the hours are contiguous, commas if not). |
| Minutes | <p>Specifies the time of day (minutes) that you want the job to be run, in relation to the specified hours:</p> <ul style="list-style-type: none">• * (Default): The job will be run on all minutes (essentially the job is run continuously, once per minute).• 0-59: The job will be run on the specified minute or minutes of the hour, where 0 is the first minute of the hour and 59 is the last minute of the hour. Use a comma or a hyphen to separate multiple minutes (hyphen if the hours are contiguous, commas if not). |

► Event Handlers

If an event handler is associated with the job, it is listed here. There are two types of event handlers:

- System event handlers, for completing system-triggered tasks. See [Managing event handlers](#).
- User-defined event handlers, for running jobs via RunEvent. See [Creating event handlers for a job](#).

To add or remove event handlers, use the **Add**, **Remove Selected**, or **Clear All** commands in the **Event Handlers** group of the **Job** tab. This group is only available when you have selected the **Event Handlers** section in the left-hand side of the job.

| Item | Description |
|------------|--|
| Active | Specifies whether the event handler is active or not within the current job. If inactive, then actions that trigger the event handler will ignore this job. |
| Event Name | <p>The name of the event handler.</p> <p>Multiple jobs can have an event handler with the same name; all those jobs will be affected when the event handler is triggered.</p> |
| Execute As | <p>The user identity under which the job will be run when the event handler is triggered.</p> <ul style="list-style-type: none"> • Owner: For system-managed event handlers, the owner is the system Scheduler identity. For user-defined event handlers, the owner is the user who last saved the job. • Requester: For all event handlers, the requester is the user who caused the event handler to be triggered. |

► Notification

This section defines email notification settings for the job. For more information, see [Setting up notifications for jobs](#).

Job variables can be used in this section. For more information, see [Using job variables](#).

| Item | Description |
|------------------------|---|
| Job Notification Level | <p>Specifies when email notifications are sent for the job. Select one of the following:</p> <ul style="list-style-type: none"> • Send all email notifications (Default) • Send email notification only when the job has errors • None • Send email notification to different email addresses when the job has errors or succeeds <p>If anything other than None is selected, then you must complete the remaining fields.</p> |
| To | The email address(es) to receive the notification email. Separate multiple addresses with a semicolon. |

| Item | Description |
|--------------------|---|
| To (on error) | The email address(es) to receive the notification email when the job fails. Separate multiple addresses with a semicolon. Only applies when Send email notification to different email addresses when the job has errors or succeeds is enabled. |
| From | The email address to use as the "From" address for the notification email. |
| Subject | The subject text for the notification email. |
| Subject (on error) | The subject text for the notification email when the job fails. Only applies when Send email notification to different email addresses when the job has errors or succeeds is enabled. |
| User Message | Optional. The body text for the notification email. Text entered here will be appended to the body text generated by Scheduler. |

► Tasks

This section defines the tasks in the job. In the ribbon, task commands are available on the **Job** tab, in the **Tasks** group.

- To add a task, click **Add**.
- To change the order of tasks, select a task and then click **Move Up** or **Move Down**.
- To delete a task, select the task and then click **Remove Selected**.
- To delete all tasks, click **Clear All**.
- To copy a task, right-click the task and then click **Copy**. You can copy the task within the same job, or to another open job in the Scheduler window. Right-click any task (or the **Tasks** section header) and then select **Paste**. The job is pasted underneath the job you right-clicked (or at the end of the list if you right-clicked the **Tasks** section header).
- To rename a task, double-click the task name to make it editable, and then type the new name. For example, if you have a job with multiple File Processing tasks, then you may want to edit the name of each task so that you know which file each task relates to at a glance. (You can also right-click and select **Rename**.)

Tasks are processed in the order they are listed in the job. By default, when you add a new task to a job, it is placed at the bottom of the list. Make sure to move the new job if it should not be processed last.

Tasks can be processed concurrently instead of sequentially if they are configured to be run as a subordinate job within the parent job.

Each task type has its own unique settings in addition to the standard task settings. For more information, see [Scheduler Task Reference](#).

► Job Results

Displays historical results for the job. This section is blank if the job has never been run.

Job results may be purged periodically by using the **Job Results Cleanup** option for the job, or by the system **Purge System Data** task.

NOTE: Users with the **Scheduled Jobs User** security permission only see results for jobs that they executed. Administrators see results for all executions.

For more information on job results, see [Viewing job results](#).

Scheduler Administration

Using the Scheduler dialog, you can review the job schedule and job results, as well as perform Scheduler administration activities such as managing system jobs and configuring Scheduler servers.

Managing scheduled jobs

The **Scheduled Jobs** tab of the Scheduler dialog displays all jobs that are scheduled to be processed or are in process. This includes scheduled jobs, jobs executed manually via **Run Now**, and jobs that were triggered for execution via an event handler.

By default, this list displays when you first open Scheduler. If you have changed the view in the dialog, you can return to it by clicking the following:

- On the **Service** tab, in the **Service View** group, click **Scheduled Jobs**.

If a job has a scheduling rule with a recurring schedule, only the first scheduled execution appears in the list. For example, if you have a job that is scheduled to run once a month for a year, you will not see all twelve scheduled executions in the list—you will only see the first scheduled execution. Once that instance has been run, the scheduling rule is re-evaluated and the next scheduled execution appears in the list.

In this tab, you can:

- **Stop a scheduled execution.** If you do not want a scheduled instance to be run, right-click the item and select **Remove from Schedule**. This not only removes the scheduled instance, it also inactivates the associated scheduling rule for the job (if applicable). If you want future scheduled instances of the job to proceed, you must edit the job to re-enable the scheduling rule.
- **Abort an in-process job.** If a job is already running and you want to stop it, right-click the item and select **Abort Working Job**. This will immediately abort the job regardless of what it is currently doing. Keep in mind that depending on the type of job and its stage in the process, this may result in side-effects such as files remaining locked or incomplete processing. No notifications will be sent for the aborted job. Generally, this action should be reserved for situations where a job has gotten "stuck" or was started in error.

- **Reschedule an execution.** To change the scheduled start date/time of a scheduled instance, right-click the item and then select **Reschedule Start Time**. The **Start Time** field in the grid is now editable, and you can type in a new date and/or time.
- **Refresh the list.** On the **Service** tab, in the **Actions** group, click **Refresh**. New scheduled jobs will be added to the list, and **Status** and **Due In** will be updated appropriately.

NOTE: Users with the **Scheduled Jobs User** security permission can only modify jobs that they placed on the schedule. Jobs scheduled by other users are visible, but are grayed out and unavailable for editing. Administrators can view and edit all scheduled jobs.

You can also view the job schedule and perform limited management tasks using the Scheduler area of the Web Client. For more information, see [Managing the job schedule in the Web Client](#).

► Scheduled job information

The following information displays for each scheduled job:

| Item | Description |
|--------|---|
| ID | The system-generated ID for the job execution. Each scheduled execution of a job has a unique ID. Job results are listed by execution ID. |
| Job | The name of the job. |
| User | The user identity that the job will be run as. If the job is a system job, the user is System . This is typically the name of the user who placed the job on the schedule, but not always (for example, when using an event handler that is set to run as the job owner instead of as the requester). |
| Status | Job status is either Pending (waiting to be executed) or Working (currently being executed). |
| Server | If a job is currently Working , then the server executing the job is listed here. Otherwise, this column is blank. |

| Item | Description |
|------------|---|
| Priority | <p>The priority category for the job:</p> <ol style="list-style-type: none"> 1. Manual: The job was executed manually. 2. Event Handler: The job was executed by a Scheduler event handler. 3. Scheduled Job: The scheduled instance of the job results from an active scheduling rule. 4. Subordinate Job: The job was generated as a subordinate job, from a currently executing job. <p>The priority category determines how jobs are evaluated for processing order, in conjunction with the job's Priority Elevation setting. Manual jobs are highest priority, and subordinate jobs are lowest priority. For more information, see Processing priority for scheduled jobs.</p> |
| Start Time | <p>The start time of the job. The job is eligible for immediate execution if the start time is now or passed. Jobs may not be executed right at the start time if no Scheduler threads are currently available to execute the job, or if other eligible jobs have higher priority.</p> <p>If the job is on the schedule due to a scheduling rule, the start time is based on the scheduling rule. If the job was manually executed via Run Now or triggered by an event handler, the start time is the time the execution was initiated.</p> |
| Due In | <p>The length of time until the job is due to be processed. For example, if the job is scheduled to run at noon and it is currently 11:50 AM, then the job is due to be run in 10 minutes.</p> <p>This column is intended to make it easy to see when a job will be run, without needing to calculate it based on the start time.</p> |

Viewing job results

Once a job has been run, you can view the results in the **Scheduler** dialog. Job results can tell you:

- Whether the job processed successfully or failed
- The process steps performed by the job, if it was successful
- The error message for the job, if it failed
- The date/time when the job started processing and how long it took to process
- The Scheduler server that processed the job

NOTE: Users with the **Scheduled Jobs User** security permission can only see job results for jobs that they executed. Administrators can see job results for all jobs.

To view results for all jobs:

- On the **Service** tab, in the **Service View** group, click **Job Results**.

The **Result History** tab opens, listing a summary of results for all jobs that have started or finished. If you want to see more specific details for a particular job execution, double-click it. This opens the related job to the **Job Results** section, where you can view more details such as the specific error message for a failed execution.

TIP: Alternatively, you can open a job directly, and go to the **Job Results** section to view results for that job only.

You can also view job results using the Scheduler area of the Web Client. For more information, see [Viewing job results in the Web Client](#).

► System job results

By default, system job results are hidden in the **Result History** tab. System jobs such as the SMTP message delivery job run frequently, and can easily fill up the result history screen, making it difficult to find results for user-initiated jobs.

If you want to view results for system jobs, you can do one of the following:

- At the bottom of the **Result History** tab, clear the **Hide system jobs** check box.
- Open the system job directly, and view its job results within the job. For example, you can go to the **Scheduled Jobs** tab and double-click the **System.SystemDataPurge** job to view all results for that job.

► Purging job results

Over time, the job result history can get quite large, so it is recommended to purge the result history periodically. There are two ways to purge result history:

- Each job can be configured to purge its own prior result history when it is run (**Job Results Cleanup**). This setting is located in the **General** section of the job properties. For more information, see [Job properties](#).
- The **System.PurgeSystemData** system job purges result history whenever it is run (by default, once per hour). For more information, see [Configuring Scheduler system jobs](#).

When you purge job results, you specify a number of days of history to be kept. For example, you may always want to keep 5 days' worth of job history, and purge any results older than that.

Managing Scheduler servers

On the **Servers** tab of the Scheduler dialog (**Service > Servers**), you can view the status of all Scheduler servers for the current system, enable or disable the servers, and change server configuration settings.

NOTE: Only administrators can manage Scheduler servers. Users with the **Scheduled Jobs User** permission cannot see the **Servers** tab in the Scheduler dialog.

This tab lists all Scheduler servers for this Axiom installation, including system Scheduler Services for on-premise systems and Cloud Integration Services for Axiom Cloud systems. For each server, the following properties are listed:

| Item | Description |
|-------------------------|--|
| Server | The server name, including the following appended codes where appropriate: <ul style="list-style-type: none">• -System indicates the system Scheduler Service• -CIS indicates a Cloud Integration Service |
| Version | The Axiom version installed on the server. This should match the version installed on the application server. |
| Service Status | The status of the service on the server, such as stopped or started. |
| Processor Status | The status of the server, such as running or stopped. |
| Processor Enabled | Whether the server is enabled for processing. |
| Has Embedded App Server | Whether the server uses an embedded application server. This is determined during the installation. |

This tab can be useful in troubleshooting—for example, if scheduled jobs are not executing, you can check this tab to see if the Scheduler service has been stopped on a server or if the server is down.

► Enabling or disabling server processing

When a new Scheduler server is installed, you must manually enable it within the Scheduler dialog before it will begin processing jobs. Later, if you want to temporarily prevent a Scheduler server from processing jobs, you can disable processing for that server.

NOTE: The system Scheduler service for on-premise systems is enabled automatically. You do not need to manually enable this service.

To enable or disable server processing:

- Select the server in the list. In the **Configuration details** section for that server, select or clear the **Schedule Processing Enabled** check box, and then click **Update**.

If a server is disabled, it will not process any new tasks. If the server is currently processing a job when it is disabled, it will finish the current task and then stop processing. The job's settings determine whether the stopped job is rescheduled for processing.

► Configuring Scheduler server settings

Each Scheduler server has its own configuration settings that determine behavior such as how many threads are available to process jobs at any one time.

NOTE: We recommend leaving the default Scheduler server configuration settings unless Axiom Support has instructed you to change them, based on your particular system and server environment.

To view or edit the configuration settings for a server, select the server in the list. The settings for that server display in the **Configuration details** section. If you edit a setting, click **Update** to apply the change.

| Item | Description |
|--------------------|--|
| Max Worker Threads | <p>The maximum number of threads that the Scheduler server has available to process jobs simultaneously. The default is 4.</p> <p>As a guideline, the Scheduler service can typically handle 2 threads per CPU on the Scheduler server. For example, a 4 CPU server could be set to 8 threads. Using a greater number of threads may degrade performance.</p> <p>NOTES:</p> <ul style="list-style-type: none">• This setting is inapplicable to the system Scheduler service. It creates threads as needed to process system jobs.• A Scheduler server will temporarily exceed the configured number of threads if a job's priority elevation is set to Interrupt and no threads are currently available to process it. |
| Server Poll Rate | <p>The rate at which the Scheduler server checks for new jobs to process, in seconds. The default is 20 seconds.</p> |
| Retry Rate | <p>The rate at which the Scheduler server will resubmit commands to the application server, in the event the original connection failed. The default is 5 seconds.</p> |

► Removing server registration

If a Scheduler server is no longer valid, or if the Scheduler software was installed on a server in error, you can remove the server registration for Scheduler.

To remove the registration for a Scheduler server:

- Right-click the server and then select **Remove Server Registration**.

The server is removed from the list. Registration can only be removed if no historical job results exist for the server.

If the service is restarted on the machine and the Scheduler installation is pointed to this Axiom installation, then the server will be re-registered.

Managing event handlers

Event handlers are used to trigger Scheduler jobs based on an event. Scheduler has a number of built-in system-managed event handlers, to support certain features. Scheduler also supports user-defined event handlers.

In the **Scheduler** dialog, on the **Event Handlers** tab (**Service > Event Handlers**), you can view all active event handlers. This includes system-managed event handlers and user-defined event handlers. If an event handler in a job is inactive, it is not listed here.

NOTES:

- User-defined event handlers cannot be created in this tab; they can only be created within a Scheduler job. For more information, see [Creating event handlers for a job](#).
- The **Event Handlers** tab is only visible to administrators. Non-admin users with the **Scheduled Job Users** permission can only manage event handlers within the jobs they have access to.

For each event handler, the following information is displayed.

- **Event Name:** The name of the event handler. If the event handler is associated with a file group, then the file group name displays in parentheses after the event name. When using the RunEvent command, it is possible to specify a file group context to target event handlers that are associated with a specific file group.
- **Job:** The job associated with the event handler. For system event handlers, the job is a system job. For user-defined event handlers, the job is the job where the event handler was created. Multiple jobs can contain a user-defined event handler with the same name (in which case the event handler is listed multiple times).
- **User:** The user identity that will be used to execute jobs triggered by the event handler. If the event handler is configured to run as the requester, then (**Requester**) is listed here. If the event handler is configured to run as the owner, then the owner name is listed here (either a specific user name, or **System**).

If you want to change how the event handler runs, you can double-click the event handler name to open and edit the associated job. Go to the **Event Handlers** section and then edit the **Execute As** setting for the event handler.

► Removing event handlers

You can remove an event handler by right-clicking the event handler and selecting **Remove event handler**.

IMPORTANT: You should only remove a system event handler if the event handler has become corrupted and you need to restore a new copy of it. If you remove a system event handler and then do not restore a new copy, the associated system functionality will become unavailable.

For user-defined event handlers, removing the event handler inactivates it within the job; it does not delete the event handler. If you want to delete the event handler entirely, you must edit the job to delete it.

► Creating system jobs and system event handlers

If any system jobs or event handlers are accidentally or deliberately removed from the system, you can restore new copies by right-clicking any event handler and selecting **Create system jobs and event handlers**. Any missing event handlers or system jobs will be re-created. You must refresh the Event Handler tab to see the re-created event handlers.

► System event handlers

Axiom uses the following system event handlers.

| Event Handler | Description |
|----------------------------------|--|
| System.ApplyUserTableConstraints | Applies constraints to tables as needed. Operates as a system job. |
| System.RebuildPartitionedTables | Rebuilds partitioned tables as needed when the partitions change. Operates as a system job. |
| System.RefreshDocumentList | Processes individual documents as part of various utilities. The job is performed as the user who executed the utility. The job is not processed on the system Scheduler server. |
| System.RefreshPlanFiles | Creates a Process Plan Files job when a user selects to run the Process Plan Files utility (from the file group menu) on the Axiom server instead of locally. The job is performed as the user who executed the utility. The job is not processed on the system Scheduler server. |
| System.SMTPMessageDelivery | Sends email notifications for Axiom. Operates as a system job. |

Managing remote data connections

Remote data connections provide a way for the Axiom Cloud to communicate with the client's local network environment. This connection can be used to support the following features:

- Importing data to the cloud from a local file or database

- Exporting data from the cloud to a local database.
- Authenticating users for access into Axiom (Windows Authentication)
- Saving output files for file processing to a local file system (snapshot or export)

For a full-process discussion of how to set up and use remote data connections with the Axiom Cloud, see the separate *Axiom Cloud Technical Guide*.

Remote data connections are managed in the **Remote Data Connections** screen of the **Scheduler** dialog. The following is required to create a remote data connection:

- The Axiom Cloud Integration Service must be already installed on a server within your local network environment, and the server must be enabled for use (on the **Server** screen of the Scheduler dialog).
- You must have the connection string to your system's namespace on the Axiom Cloud. If you do not have this information, contact Axiom Support.

NOTE: Only administrators can create and edit remote data connections. Users with the **Scheduled Jobs User** permission cannot see the **Remote Data Connections** tab in the Scheduler dialog.

To create a remote data connection:

1. On the **Service** tab of the Scheduler dialog, click **Remote Data Connections**.
2. At the bottom of the dialog, click **New Connection**.

A set of empty fields displays in the **Configuration details** section.

3. Complete the following settings:

| Item | Description |
|--------------------------------|---|
| Name | The name of the remote data connection. |
| Description | Optional. The description of the remote data connection. |
| Connection String (SAS) | The connection string to the namespace that was created by Axiom Support in the Microsoft Azure environment for your Axiom Cloud system. If you do not have this connection string, contact Axiom Support. |
| Use for authentication service | Specifies whether the remote data connection is used for Windows Authentication. This only applies if Windows Authentication is enabled for your Axiom Cloud system. If enabled, this remote data connection will also be used for any Active Directory Import tasks set up for your system. |
| Cloud Integration Server | The name of the Axiom Cloud Integration Service installed in the local network environment. The name of this server is ServerName-CIS . |

4. Click **Save Connection**.

The connection now displays in the list of remote data connections. You can now use this connection when setting up import and export utilities.

► Editing or deleting remote data connections

- To edit a connection, select the connection in the list and then make the necessary edits in the **Configuration details**. Click **Save Connection** to save the changes.
- To delete a connection that you no longer need, select the connection in the list and then click **Delete Connection**.

► Accessible folders

When the Axiom Cloud Integration Service is installed, you can optionally specify a list of accessible folders for the service. The service is then limited to using those folders when reading files or saving files. If no accessible folders are specified, then the service can use any folder that it has permission to access (based on the account the service is running as, which by default is the local service account).

You can view this list of folders in the remote data connection properties, so that you know which folders are accessible to the connection. To view the accessible folders, select the connection in the list and then view the **Accessible Folders** field at the bottom of the **Configuration details**. This read-only field displays either a list of folders or the text **All (no folders configured)**.

The list of accessible folders cannot be defined or modified within Scheduler. If you need to change the list of accessible folders, you must use the Axiom Software Manager to repair the Axiom Cloud Integration Service. For more information, see the *Axiom Cloud Technical Guide*.

► Default remote data connection

Axiom automatically specifies a default remote data connection for the system. The name of this connection can be returned as follows:

- In import utilities, you can use the system variable `{DefaultRemoteDataConnection}`.
- In a spreadsheet, you can use the function `GetSystemInfo("DefaultRemoteDataConnection")`.

If your system only has one defined remote data connection, then that connection is the default connection. If your system has multiple remote data connections, then the default connection is selected alphabetically from all connections that are not enabled for authentication. If all connections are enabled for authentication, the selection is simply alphabetical.

Scheduler Task Reference

Each Scheduler task in a job has two sections of task properties:

- **Task Control:** Common task properties that apply to all task types. For more information, see [Task Control properties](#).
- **Task Details:** Properties specific to the current task type. For more information, see the topics for each individual task type.

The following task types are available:

| Task | Description |
|---|--|
| Active Directory Import | Import users from Active Directory into Axiom. This task adds new users, and can also disable users that no longer exist in the Active Directory domain. |
| Collect Worksheets | Collect worksheets from multiple files into a single file. |
| Copy On Demand Plan Files | Copy plan files from one on-demand file group to another. |
| Create Plan Files | Create new plan files (same as the Create Plan Files utility for file groups). |
| Echo Task | Test the Scheduler server. This task sends a message to the Scheduler server and asks it to send the message back. |
| Execute Command Adapter | Execute a command from the Command Library. |
| Execute SQL Command | Run a SQL statement on an Axiom database. |
| Export ETL Package | Export data to an external database, using an export utility defined in the Exports Library. |
| File Processing | Perform file processing actions on a report. You can use the report's native file processing settings, or override the settings. |
| Import ETL Package | Import data into Axiom, using an import utility defined in the Imports Library. |

| Task | Description |
|--|--|
| Process Document List | Process any set of Axiom files—for example, driver files or report utilities. The task calculates and saves the files, and can also refresh Axiom queries and save data to the database. |
| Process Plan Files | Process plan files (same as the Process Plan Files utility for file groups). |
| Process Template List | Process a template file. The task runs designated Axiom queries, time-stamps the queries, and saves the template. |
| Purge System Data | Purge old Scheduler results and system temp tables. |
| Raise Event | Trigger another Scheduler job for execution, using a named event handler. |
| Run Scheduler Job | Run another Scheduler job as a subordinate job. |
| SMTP Message Delivery | Deliver email notifications resulting from Scheduler jobs. |
| Start Process | Start a process definition for Process Management. |
| Update Indexes and Constraints | Update the indexes and constraints in your Axiom database. |

Task Control properties

The following task properties are available for all Scheduler task types. To edit these properties, select the task in the Scheduler job, then expand the **Task Control** section.

| Item | Description |
|---|---|
| Skip this Task | <p>If selected, the task will not be run when the job is processed.</p> <p>By default, this option is not selected, which means this task will be run.</p> |
| Process task only if the value of this expression is true | <p>Optional. Enter a logical expression to conditionally process this task depending on whether the expression resolves to true or false at the time the job is executed. If true, the task is processed as normal. If false, the task is skipped.</p> <p>The logical expression is evaluated by the Scheduler server using an IF function. The expression can be any statement that would be valid within an IF function. You can use Excel functions, Axiom functions, and Scheduler job variables in the expression. If you use a job variable in the expression, you must place the variable in double quotation marks unless you expect it to be resolved and evaluated as a number.</p> <p>For more information, see Conditionally processing tasks in a job.</p> |

| Item | Description |
|--|--|
| If this Task fails, continue executing subsequent Tasks | <p>If selected, the job will continue processing even if this task fails.</p> <p>By default, this option is not selected. If a task in a job fails, the job is canceled and no further tasks are processed.</p> |
| Create a Subordinate Job for this Task | <p>If selected, this task will be processed as a subordinate job to the current job.</p> <p>Selecting this check box enables concurrent processing of different tasks, if the option to Wait for all Subordinate Jobs to complete before proceeding to the next Task is not selected.</p> <p>NOTE: This option is not available for Process Plan File tasks.</p> |
| Wait for all Subordinate Jobs to complete before proceeding to the next Task | <p>If selected, the job will wait for any subordinate jobs to complete before moving to the next task.</p> <p>If this check box is not selected, and the option Create a Subordinate Job for this Task is selected, then tasks can be processed concurrently instead of sequentially.</p> <p>This check box is selected by default for Plan File Refresh and File Processing tasks. For other task types, this option is not selected by default.</p> |
| Override Log Level for this Task | <p>By default, Scheduler jobs perform logging at the same level that is specified for the application. If necessary, you can override the logging level for a particular task, so that it always runs at a specified logging level. You may want to do this if you encounter performance issues for tasks that generate a lot of logged messages.</p> <p>To do this, select the check box for Override Log Level for this Task, then select the desired logging level from the drop-down list.</p> <p>NOTE: This option is only available for File Processing tasks.</p> |

NOTE: Older systems may see a setting named **Workbook processing engine to use**. If this option is present, it should always be set to **Axiom Web Engine**. Use of Excel for processing on the Scheduler server is no longer supported. All Scheduler tasks that involve spreadsheet processing are processed using the same spreadsheet emulation engine as the Windows Client.

Iteration

This section can be used to optionally enable iterative processing for the task. For more information, see [Using iterative task processing](#).

| Item | Description |
|---|---|
| Iterate this Task | Specifies whether iterative processing is enabled for the task. If enabled, then the task will be performed N times, where N is the number of unique items in the specified iteration column. Job variables can be used to apply the current iteration value and iteration number to the task. |
| Create a Subordinate Job for each iteration | <p>Specifies whether each iteration is processed as a separate subordinate job. By default, this is disabled, which means that all iterations are processed sequentially within the overall subordinate job created to process the iterations.</p> <p>If enabled, then each iteration is processed as a separate subordinate job, enabling concurrent execution of multiple iterations. This option should only be enabled if the order of iteration processing is not important.</p> |
| Column | <p>The column that contains the values to iterate over. Use <code>Table.Column</code> syntax to specify the column. Multiple-level lookups can be used.</p> <p>For example, if you specify <code>Dept.Region</code>, then the task will be processed once for each unique region value in the column (after applying any filter to limit the list of values).</p> |
| Group By | <p>Optional. By default, the group by column is the same as the iteration column, so that the task is processed once for each unique value in the iteration column. However, if needed, you can specify a different grouping level.</p> <p>You can use any column or columns that would be valid as the "sum by" level for an Axiom query, where the primary table is the table specified for the iteration column.</p> |
| Order By | <p>Optional. By default, the values are sorted based on the iteration column, in ascending order. You can specify a different sort column, or use the same sort column but change the order to descending.</p> <p>The sort order is ascending unless the keyword <code>desc</code> is used to specify descending order. For example:</p> <pre>Dept.Dept desc</pre> |
| Filter | Optional. A filter criteria statement to limit the list of values for the iterative processing. You can use any filter that is valid against the source table (the table of the iteration column). |

Active Directory Import task

This task imports users from Active Directory groups into Axiom security. For more information on using Active Directory integration with Axiom, see the *Security Guide*.

This task has three tabs of settings: **Source Directory**, **Notification**, and **Preview Import**.

NOTE: The user running this task must be an administrator or have the **Administer Security** permission.

For Axiom Cloud systems, the Active Directory Import task can import users from your local Active Directory by use of the Axiom Cloud Integration Service. If you have a remote data connection that is enabled for user authentication, this task will use that connection when the job is executed by Scheduler.

► **Source Directory tab**

On this tab, you specify the domain to import from and the groups to import.

| Item | Description |
|------------------|--|
| Domain or Server | <p>Select either Domain or Server to specify the source domain for the import.</p> <ul style="list-style-type: none">• If you select Domain, enter the name of the domain.• If you select Server, enter the name of the domain controller server. <p>The server option is available in case you are not currently logged into the source domain, and your current domain does not have access to the source domain. In this case, you must use domain credentials in order to access the source domain.</p> <p>Only one domain can be selected per import task. If you want to import users from multiple domains into an Axiom system, then you must create multiple import tasks.</p> |
| Credentials | <p>Specifies the credentials to use when accessing Active Directory for the import. Select one of the following:</p> <ul style="list-style-type: none">• Use process credentials: (Default) Use the credentials of the network service account for Axiom Scheduler Server (on-premise installations) or Axiom Cloud Integration Service (Axiom Cloud systems).• Specify domain credentials: Enter the credentials of a specified domain User and Password. This option is required if you identified the source domain using the server name instead of the domain name. |

| Item | Description |
|--------------------|--|
| Never Enable Users | <p>Specifies whether the import enables imported users as part of the process:</p> <ul style="list-style-type: none"> • If unchecked (default), then newly imported users are enabled as part of the import. Additionally, any existing imported users who have been changed to disabled are re-enabled. • If checked, then newly imported users are not enabled as part of the import. A security administrator must modify the security settings after the import is complete to enable the new users. Existing imported users retain their current enabled status. |
| Groups to import | <p>The Active Directory groups for which members will be imported into Axiom Security.</p> <ul style="list-style-type: none"> • Click Add to select from a list of groups for the specified domain. If the specified domain name is not valid or if Axiom cannot connect to it, then an error will result when attempting to add groups. • If you need to remove a group, select the group and click Remove. • Click Role Mapping to define mappings for the selected groups. If a mapping exists for a group, then when users are imported for that group they are automatically assigned to the mapped role and subsystem. See the discussion following this table for more information. |

Role mapping

In the **Role Mapping** dialog, click **Add mapping** (the plus icon) to add a role mapping for a group. Then complete the following:

- In the **Directory Group** column, select the Active Directory group to be mapped.
- In the **Axiom Role** column, select the role to be assigned to users in that group. If you want to map the group to more than one role, add another mapping row. You can select **None** if you do not want the users in the group to be assigned to any role.
- In the **Subsystem** column, select the subsystem for users in that group. If you want to map the group to more than one subsystem, add another mapping row. This option only displays if subsystems are enabled for your system.
- In the **User Type** column, select the license type for the imported users. The default license type is **Standard**.
- In the **Authentication Type** column, select the authentication type for the imported users, **Windows User** or **SAML**. The default authentication type is Windows User. Note that the selected authentication type will be assigned to users regardless of whether that authentication type is currently enabled for the system.

You can map each group to multiple roles and subsystems. If a group has no defined mappings, then the users will not be assigned to any roles or subsystems. If the import creates new users without mappings, the assigned user type is Standard and the assigned authentication type is Windows User.

To remove a mapping, select the mapping in the grid and then click **Remove mapping** (the X icon). If users have already been imported using this mapping, removing the mapping will not remove the users from the role or subsystem in subsequent imports (unless other group mappings in the import use the same role or subsystem, and the users are not also part of that group).

NOTE: If a user belongs to multiple mappings—either multiple mappings for a single group, or multiple mapped groups—then the user will be assigned to the user type and the authentication type for the last-processed mapping.

► Notification tab

On this tab, you specify users to be notified when changes are made in Axiom Security due to the import.

Type in one or more email addresses to be notified. Separate multiple addresses with a semi-colon. For example:

```
jdoe@axiomepm.com;jsmith@axiomepm.com
```

When the import task is run, if any users are created or modified in the Axiom system, an email notification will be sent to the addresses specified here. The email summarizes the changes made. This email notification is independent of any job-level notification settings (which notify based on overall job completion or failure).

We recommend setting up this task-level notification to send emails to the security administrator(s) responsible for maintaining the security settings in Axiom, so that he or she can define security settings for newly added users, validate changes made to existing users, and perform any other follow-up tasks.

Scheduler [job variables](#) can be used in this setting.

► Preview Import tab

On this tab, you can preview the import results to test that the import is set up as desired.

To preview the results, click **Preview**. Axiom processes the import task but does not actually make the changes to the system. Instead, the tab displays a summary of the changes that would result.

The preview shows a list of users that would be added, changed, or disabled.

NOTE: The preview is always executed locally, even for Axiom Cloud systems. The remote data connection to the Cloud Integration Service is only used when the task is executed by Scheduler.

Collect Worksheets task

This task collects sheets from multiple source workbooks and combines them into a single target workbook. You can then save the target workbook to a specified file location, and/or email the workbook.

NOTE: This task is primarily intended for backward-compatibility only. The main method of performing a file collect operation is to use the file processing feature with the [File Processing Scheduler](#) task. For more information on setting up a file collect report using file processing, see the *Axiom File Setup Guide*.

Typically, this task would be used at the end of a job with multiple File Processing tasks, to take the results of those tasks and collect them into a single workbook.

| Item | Description |
|------------------------|--|
| Save or Email Workbook | <p>Specifies the delivery option for the target workbook. Select one of the following:</p> <ul style="list-style-type: none">• Save Workbook: The target workbook is saved to the specified output folder.• Email Workbook: The target workbook is emailed to the specified recipients. The file is not saved anywhere on the file system.• Save and Email Workbook: The target workbook is both saved and emailed. |

► Target Workbook

Complete the following settings to define the target workbook:

| Item | Description |
|------------------|--|
| Output Folder | <p>The folder location where the target workbook will be saved (if you are saving the workbook). Click the folder icon to select a folder location, or type a folder location.</p> <p>If the specified folder does not already exist, Axiom attempts to create it.</p> <p>Job variables can be used in this setting.</p> |
| Output File Name | <p>The name of the target workbook. Job variables can be used in this setting.</p> |
| File Type | <p>The file type of the target workbook. Select XLS, XLSX, or XLSM.</p> <p>NOTE: PDF displays as an option, but it is not supported in this context.</p> |

► Email Settings




This section only applies if you are emailing the target workbook. The "From" address is always the Scheduler default From address (as defined in the system configuration settings).

| Item | Description |
|--------------|---|
| To | Enter the email addresses to receive the target workbook via email. Separate multiple addresses with a semicolon. |
| Subject Line | The subject line for the email. |
| Body Text | The body text for the email. |

► Source Workbooks

In this section, you specify one or more source workbooks from which to collect worksheets. Workbooks are identified by folder location. Within a folder location, you can specify one or more workbooks by name, or by using wildcards, or by using *.* to collect all workbooks at the location.

All sheets in each source workbook will be collected. Ideally, you will be collecting from workbooks that only contain relevant sheets (for example, no blank "Sheet2," etc.), and where the sheets have unique names. If multiple workbooks have sheets with the same name, the sheets will be incremented by number in the target workbook.

-  To add a workbook, click the Add button. In the **Edit Workbooks Source** dialog, complete the settings as described below, then click **OK** to add the workbook to the list.
-  To remove a workbook, select the workbook in the list and then click the Remove button. Only one workbook can be selected at a time.
-  To change the order of workbooks, select the workbook in the list and then click the arrow buttons to move the workbook up or down. Source workbooks are processed in the order they are listed in the grid.

| Item | Description |
|-------------|---|
| Folder Path | <p>The folder location of the source workbook(s). Click the folder icon to select a folder location.</p> <p>NOTE: The Folder Path location must be accessible by the Scheduler service user account. If you specify a network folder location using the Browse button, the location is automatically entered as a UNC path. If you specify a C: drive location, that will be evaluated as the C: drive of the Scheduler server.</p> <p>Job variables can be used in this setting.</p> |

| Item | Description |
|-----------|--|
| Workbooks | <p>The workbooks from which you want to collect worksheets, within the specified folder path.</p> <ul style="list-style-type: none"> Specify *.* if you want to collect all files in the folder path. Specify individual file names to collect from specific files. Separate multiple file names with semicolons. <p>You can use wildcards (* or ?) to specify groups of files that share naming conventions. For example: <code>North*.xls</code> to collect all XLS files where the file name starts with "North".</p> <p>Only files with the following file types are valid to be collected: XLS, XLSX, XLSM. If you are using wildcards, the matches must be valid file types, or else the task will fail with an error.</p> <p>Job variables can be used in this setting.</p> |

Once you have saved a source workbook location, you can edit it by double-clicking the row.

Copy On Demand Plan Files task

This task copies on demand plan files from one file group to another. It performs the same actions as the **Copy On Demand Plan Files** command in the Command Library.

This is an advanced feature and should only be used if it is the only way to achieve the desired population of plan files between two related file groups. It is the responsibility of the solution designer to ensure that the copied plan files will behave as expected in the target file group. For example, the plan file must be designed to dynamically save to the appropriate tables and columns within the context of the new file group.

The Copy On Demand Plan Files task uses two tabs to define the properties of the task.

- **Options:** Defines the options to be used for the copy operation
- **Plan Files:** Specifies the plan files to copy

► Options tab

The following options are available on the Options tab. Note that all of these options can be changed dynamically by using [system variables](#).

| Item | Description |
|---------------------------------|---|
| Source File Group | The file group to copy plan files from. Click the folder icon to select a file group. You can select any on-demand file group, or any file group alias that currently points to an on-demand file group. |
| Destination File Group | The file group to copy plan files to. Click the folder icon to select a file group. You can select any on-demand file group, or any file group alias that currently points to an on-demand file group. |
| Keep original plan file creator | <p>Specifies whether the plan file creator for the copied plan files is set to the same creator as the original plan files. By default, this option is enabled.</p> <p>If this option is disabled, then the plan file creator for the copied plan files is set to the user identity used by the Scheduler job when it is run.</p> |
| Use default template | <p>Specifies whether the copied plan files have the option to adopt the default template of the new file group. This is primarily intended to be used when copying plan files to a file group that uses virtual, form-enabled plan files, so that the copied plan files can be converted to virtual files and use the new template.</p> <ul style="list-style-type: none"> • If disabled (default), then the target file group must contain copies of the original templates that were used to create the plan files from the source file group. If these templates are not present, then the copy process will fail. • If enabled, then the copied plan files will be assigned a template as follows: <ul style="list-style-type: none"> ◦ If the target file group contains copies of the original templates that were used to create the plan files from the source file group, the copied plan files use those templates. ◦ If the target file group does not contain copies of the original templates, the copied plan files use the default template specified for the target file group in the file group properties. <p>If the target file group does not contain copies of the original templates and does not have a designated default template, then the copy process will fail.</p> |
| Copy plan file attachments | <p>Specifies whether plan file attachments are copied to the target file group when a plan file is copied. By default, this option is enabled.</p> <p>If this option is disabled, then plan file attachments will not be copied to the target file group.</p> |

| Item | Description |
|----------------------------|--|
| Save plan files after copy | <p>Specifies whether the new plan files are processed and saved in the target file group after the copy is performed. This is intended to perform a save-to-database within the context of the new file group. By default, this option is disabled.</p> <p>If you enable this option, then after the plan files are copied to the new file group, they are opened, refreshed, and saved (including a save-to-database). The refresh includes all active Axiom queries where Refresh during document processing is enabled.</p> <p>Regardless of whether this option is enabled, if it is ever intended to save the copied plan files in the target file group, then they must be designed so that they save data to the appropriate tables after being copied.</p> <p>NOTES:</p> <ul style="list-style-type: none"> • If Process with Utilities is enabled for the target file group, then utility processing is performed instead of normal processing. The default data source is used. • If you enable this option but also specify a Copy data utility, then the new plan files are not processed and saved. Instead, the designated utility file is processed for each new plan file. |
| Copy data utility | <p>Optional. Specifies a utility file to process for each copied plan file. You can select any file in the Utilities folder of the target file group, or a file in the Reports Library.</p> <p>The primary purpose of this option is to handle copying virtual plan files between file groups. Because the plan files are virtual, no data exists in the file itself and therefore saving the new plan file will not populate data for the new file group. Instead, you should create a utility file that queries in the necessary data for the original plan file, then saves the necessary data for the new plan file to the appropriate tables for the new file group. Reserved document variables are available to return information in the utility file such as the old plan file code and the new plan file code.</p> <p>For more information, see Copy data utility.</p> <p>NOTE: Save plan files after copy must be enabled in order to specify a copy data utility. If a utility is specified, then the new plan files are not saved and instead the utility file is processed for each new plan file.</p> |
| Default Values | <p>Optional. This section can be used to apply default values to any columns in the target plan code table, when the new record is created in the target file group. For more information, see Defining default values.</p> |

► Plan Files tab

On the **Plan Files** tab, specify the plan files from the source file group that you want to copy to the target file group. There are three different options that you can use to specify the plan files: **Choose from list**, **Use filter**, and **All**.

The most common option when copying plan files using Scheduler is to define a filter. You can dynamically copy a subset of designated plan files using the filter. If the Scheduler task is triggered by using RunEvent, you can pass in the filter from the source of the RunEvent (such as an Axiom form).

Copy a filtered set of plan files

To use a filter to copy a subset of plan files, select **Use Filter**. When the Scheduler task is executed, Axiom will process only the plan files that meet the filter. You can specify the filter directly, or use a job variable.

To specify the filter, click the Filter Wizard button. You can also manually type a filter criteria statement into the filter box. The filter must use the plan code table of the source file group, or a lookup table. For example: `CapReq2020.Transfer=1`, where CapReq is the plan code table.

Once you have entered a filter, you can click **Refresh plan file list** to show the plan files that currently match the filter. The refresh feature is intended to help you determine whether you have defined the filter correctly.

If you want to set the filter dynamically, you can use the Filter [system variable](#) to override the filter defined in the task. This is intended for use when running Scheduler jobs via RunEvent. If a variable value is specified when triggering the event, such as the value `CapReq2020.CapReq IN (45, 67, 98)`, then that filter statement is used to determine the plan files to be copied instead of the filter defined in the task.

Copy all plan files

To copy all plan files, select **All**. When the Scheduler task is executed, Axiom will copy all plan files in the file group (except for those hidden via the Show on List column). This is not a common use case for the copy feature, but can be used if needed.

Copy selected plan files

To copy certain plan files, select **Choose from list**, and then select the check boxes for the plan files that you want to copy. When the Scheduler task is executed, Axiom will copy only the selected plan files. This is not a common use case for the copy feature, but can be used if needed.

NOTE: This option is not available when using a file group alias as the source file group for the task. This is because the alias could change to point to any file group, which could result in a different list of plan files.

► Defining default values

When the copy action is performed, the columns for the new record are populated as follows:

- If a value has been defined for a column in the **Default Values** section, that value is used.
- Otherwise, the value from the original record in the source file group is used. This only occurs if the column names match in the source and target tables, and if the column in the target table is a compatible data type to accept the copied value.

If a column exists in the source table but not the target table, that value is ignored and does not cause an error. If a column exists in the target table but not in the source table, then it is only populated during the copy action if a default value has been defined. If the target table contains columns with lookup relationships, those columns must be populated with valid values (either from the original record or by using default values) or else the copy action will fail.

To define default values for the new records:

- Click the plus button **+** to add a new column/value pair to the **Default Values** section.
- In the left-hand box, type the name of the column in the target plan code table. For example: `SourceID`. Do not use `Table.Column` syntax.
- In the right-hand box, type the value to be placed in this column. You can enter a "hard-coded" value, or you can enter the name of a column from the source plan code table in brackets to use the value from that column. For example, `[CapID]`. The column reference is only necessary if you want the source column value to be placed in a column that has a different name than the source column. If the columns have the same name, the value will be copied automatically as noted previously in this section.

For both the column name and the value, you can use file group variables via a file group alias. Axiom looks up the current target of the alias, and finds the current value of the designated variable within that file group. Built-in variables and custom variables can both be used. To reference a variable, use the following syntax:

```
{FileGroupAliasName.VariableName}
```

For example: `{CP_CurrentYear.FileGroupYear}` returns the file group year for the file group that is currently the target of the `CP_CurrentYear` alias.

Scheduler job variables can also be used in the column name and in the value.

► Overriding task settings using system variables


All of the settings for the Copy On Demand Plan Files task can be overridden using system variables. This is intended for use when the task is being triggered by `RunEvent` (such as from within an Axiom form), and you want to pass in variable values to determine how the task is run.

The variable names for this task are as follows:

| Variable | Description |
|-----------------------------|---|
| CopyDataSourceFileGroupID | Overrides the Source File Group . Must be set to a valid file group ID. File group names or alias names cannot be used. |
| CopyDataTargetFileGroupID | Overrides the Destination File Group . Must be set to a valid file group ID. File group names or alias names cannot be used. |
| CopyDataUtilityPath | Overrides the Copy data utility . Must be set to a valid document path in Axiom. |
| Filter | Overrides the Plan File Filter to specify the plan files to copy. Must be set to a valid filter criteria statement. |
| KeepOriginalPlanFileCreator | Overrides the option Keep original plan file creator . Must be set to a valid Boolean value (True/False). |
| UseDefaultTemplate | Overrides the option Use default template . Must be set to a valid Boolean value (True/False). |
| CopyPlanFileAttachments | Overrides the option Copy plan file attachments . Must be set to a valid Boolean value (True/False). |
| SavePlanFilesAfterCopy | Overrides the option Save plan files after copy . Must be set to a valid Boolean value (True/False). |

To override task properties using these variables:

- Add the variables that you want to use to the **Job Variables** tab. For example, if you want to override the source and target file groups, the copy data utility, and the plan file filter, then add those variables to the Job Variables tab. You do not need to add a variable name if you do not plan to override it.



Example Job Variables tab to override certain settings for the copy task

You do not need to define a default value for the variable. If the value is blank, then the setting defined in the task is used. The corresponding task property will only be overridden if the variable has a defined value.

- You do not need to add the variables to the task properties. The variables automatically overwrite the task properties if they have defined values.

- When configuring RunEvent, define values for the variables as needed. For example, you could have a form where you allow the user to select the source and target file group for the copy action. Based on the user's selected file group names, you can use the GetFileGroupID function to determine the IDs for those file groups. You can then pass those IDs as variable values for the variables CopyDataSourceFileGroupID and CopyDataTargetFileGroupID.

Example RunEvent properties to pass certain variable values to the copy task

► Plan file process considerations

If the target file group has an active plan file process, the new plan file is started in that process as part of the plan file creation. The process initiator for the plan file is set as follows:

- If the plan file process has a designated Process Initiator Column, the user listed in that column is the process initiator.
- If the plan file process does not have a Process Initiator Column, or the column value is blank, then:
 - If **Keep original plan file creator** is enabled for the command, then the original plan file creator is the process initiator.
 - Otherwise, the user performing the copy operation is the process initiator.

► Copy data utility

If a **Copy data utility** is specified, this processing is performed as follows:

- The selected plan files are first copied to the new file group. If the plan files are virtual, then the placeholder document records are copied instead of physical plan files.
- The utility file is opened once before processing begins. Any data lookups or Axiom queries that are configured to refresh on open are executed at that time.
- The utility file is then iteratively processed for each new plan file as follows:
 - Document variables are set in the utility, and the workbook is calculated.
 - Axiom queries set to **Refresh during document processing** are refreshed.
 - A save-to-database is executed.

The utility file is *not* closed and reopened for each new plan file. All processing occurs within the same file session, similar to when performing multipass file processing.

The following reserved document variables are available to the utility file, to be returned using `GetDocumentInfo`. These variables return necessary information about the copied plan files and the source and target file groups.

| Variable | Description |
|-------------------|---|
| SourceFileGroupID | The ID of the source file group. You can use this ID in functions such as <code>GetFileGroupVariable</code> —for example, to return the name of the data table to query from the source file group. |
| SourcePlanCode | The plan code of the original plan file from the source file group. You can use this code to filter Axiom queries to return data for the original plan file. |
| TargetFileGroupID | The ID of the target file group. You can use this ID in functions such as <code>GetFileGroupVariable</code> —for example, to return the name of the data table to save data to for the target file group. |
| TargetPlanCode | The plan code of the new plan file in the target file group. You can use this code to save data for the new plan file. |

For example, `GetDocumentInfo("Variable", "SourceFileGroupID")` returns the ID of the source file group.

Create Plan Files task

This task creates plan files for a file group. It works the same way as the **Create Plan Files** utility that is available from the file group menu.

This task has two tabs of settings in the **Task Details** area: **General** and **Plan Files**.

NOTE: If you are using Create Plan Files to create new on-demand plan files, those plan files will be automatically started in the designated **Plan File Process** for the file group. This only applies when creating a brand new plan file. If an existing plan file is overwritten, its process status will be left as is.

► General tab

The following settings are available on the General tab:

| Item | Description |
|--------------------------------|--|
| Select File Group | <p>The file group for which plan files will be created. You can select any file group or file group alias.</p> <p>If the Scheduler job is stored in a file group Utilities folder, then you can select Use Current File Group to automatically use the file group that the Scheduler job belongs to. This is the recommended approach when the Scheduler job belongs to a file group, so that it will automatically update to point to the current file group when the file group is cloned.</p> <p>NOTE: If the task uses an alias, then you cannot select individual plan files on the Plan Files tab. Only the Use Filter and All options are available.</p> |
| Overwrite existing plan files? | <p>By default, this option is not selected, which means that existing plan files will not be overwritten, even if the plan file is selected to be created.</p> <p>If selected, existing plan files will be overwritten.</p> |

► Plan Files tab

On the Plan Files tab, specify the plan files that you want to create. This tab lists all plan codes that you have the right to access. (If a plan code has been set to **False** in the designated **Show On List Column** for the plan code table, then it is not available in this list.)

You can create plan files in any of the following ways:

- **Create all plan files:** To create all plan files, select **All**. This will cause all plan files to be created, for all existing and future plan codes.

Alternatively, you can select **Choose from list** and then select the check box in the column header, causing all plan codes to be selected, but then the list of plan codes is fixed and will not adjust for any future changes. For example, if you add a new department in the future, that new department will only be created by this task if you use the **All** option.

- **Create selected plan files:** To create certain plan files, select **Choose from list** and then select the check boxes for the desired plan codes.

To find the plan files you are looking for, you can sort, filter, and group the list using standard Axiom grid features. You can show additional columns and hide columns by right-clicking in the column header. If you have filtered the list, you can select the check box in the header to select only the plan codes that currently display in the dialog.

NOTE: This option is not available if the file group for the task is an alias. This is because the list of plan files could change when the alias target changes.

- **Create a subset of plan files using a filter:** To use a filter to create a subset of plan files, select **Use filter**, and then type a filter into the filter box. You can also use the Filter Wizard to build the filter. The filter must use the plan code table or a reference table that the plan code table links to. For example: `DEPT.Region='West'`.

Once you have entered a filter, you can click **Refresh plan file list** to show only those plan codes that currently match the filter. This feature is to help you determine whether you have defined the filter as intended. The filter will be applied to the list of plan codes when the Scheduler job is processed, so if changes have been made to the plan code table since then, the actual list of plan files processed will reflect those changes.

You can also use a job variable for the filter. For example, you can define a job variable named "filter" and then place the text `{filter}` in the filter box. This is intended for use when running Scheduler jobs by using the RunEvent function. If a value is specified in the RunEvent function, such as "Filter=dept.region='west'", then that filter will be used in place of the `{filter}` variable to determine the list of plan files to be created.

NOTE: If you use a variable, and you leave the default value for that variable blank within the **Job Variables** tab, then all plan codes will be created if no value is passed by the RunEvent function (or if the value is invalid). You may want to define a default filter that results in no values (such as `1=0`), so that plan files are only created if a valid filter value is passed.

IMPORTANT: For all of these options, the **Overwrite existing plan files** option on the General tab determines whether all selected plan files are created, or only the plan files that do not already exist.

Echo task

This task is primarily used for testing purposes, to check whether a Scheduler server is running and operational. The task sends a message to the Scheduler server, and asks it to send the message back (an "echo"). If successful, the message displays in the job results. No other action is performed.

| Item | Description |
|-----------------|---|
| Message to Echo | The message to send to the Scheduler server for testing. Job variables can be used in this setting. |
| Sleep Time | The time to pause in between message echoes, in seconds. Scheduler will echo the message once, then wait the specified sleep time, then echo the message again. |

► Using an Echo task to control jobs with sub-jobs

If the **Task Control** option **Wait for all Subordinate jobs to complete before proceeding to the next task** is enabled for an Echo task, then the Echo task will inherit the status of the sub-jobs when they are completed. If one of the sub-jobs fails, then the Echo task will inherit that failed status. By default, this means that the job will not process any further tasks due to the sub-job failure. This behavior can be used to stop a job from processing tasks if previous sub-jobs fail.

For example, you could have a Scheduler job with several import tasks configured to run as sub-jobs, followed by some document processing tasks. Since the document processing tasks depend on the imported data, you do not want the document processing tasks to be run if any of the import tasks fail. However, because the import tasks are running as sub-jobs, the sub-job failure does not stop the parent job from continuing to process tasks. You can place an Echo task after the import tasks and configure it to wait, and then if any sub-jobs fail the Echo task will fail.

Execute Command Adapter task

This task executes a selected command from the Command Library.

► Task properties

This task has one property named **Command Name** that specifies the command to execute.

To select a command to execute:

1. Click **Edit Command**.
2. In the **Shortcut Properties** dialog, click the browse button [...] to the right of the **Shortcut Target** box.
3. In the **Axiom Explorer** dialog, select the desired command from the **Command Library**, then click **Open**.

This returns you to the Shortcut Properties dialog. The selected command is now listed in the Shortcut Target box, and the Shortcut Parameters section displays the parameters for the command.

4. Complete the **Shortcut Parameters** for the command as needed. The available parameters depend on the selected command.

You can later edit the shortcut parameters or select a different command by clicking **Edit Command**.

► Supported commands

Only certain commands are available for execution in this context. The following commands are available:

- File Group Rollover
- Create File Group Scenario

Systems with installed products may have Scheduler jobs that use the following additional commands:

- Create File Group From Prototype
- Upgrade File Group To Prototype Version


These commands can only be configured by product developers. Clients looking for more information on how to use a particular product-delivered Scheduler job should consult their product documentation. Syntellis employees should consult internal resources for more information on this feature as needed.

Execute SQL Command task

This task runs a SQL statement on an Axiom database. If needed, you can also use this task in a user-defined job to run any valid SQL statement on an Axiom database.

NOTE: The SQL statement in this task will be run as the Axiom database user, regardless of which user executes the job.

This task has the following settings:

| Item | Description |
|-----------------------|---|
| Source Axiom Database | Select the database on which to run the SQL statement: <ul style="list-style-type: none">• Current system database: The database for the current system.• Current audit database: The corresponding audit database for the current system. |
| SQL Command Text | <p>Enter any valid SQL statement to be run against the specified database.</p> <p>To validate the syntax of the SQL statement, click the Check SQL syntax button . Axiom sends the statement to your database server to see if the statement can be parsed, resulting in either a success message or an error message.</p> <p>Job variables can be used in the SQL statement. The Check SQL syntax button is not available if the statement uses variables.</p> |

Export ETL Package task

This task exports data from Axiom to an external database (same as executing an export from the **Exports Library**).

This task has one setting, **Select ETL Export Package**. This is the name of the export package to process. You can select any export that is defined in the current system.

File Processing task

This task performs file processing on a specified report file or file group utility. The file must already be enabled for file processing. You can use the file processing settings that are already in the file, or you can override any setting.

The following settings must be completed for the task:

Item

Description


File to Process

The report to process for the task. Click the **Browse** button to open the **Axiom Explorer** dialog, and then select a report to process. You can select any report that you have access to within the Reports Library or a file group Utilities folder.

Only one report can be selected for each File Processing task. If you want to process multiple reports, you can add multiple File Processing tasks to the Scheduler job.

TIP: Once the file is selected, only the file name displays in the task. If you want to know the folder location of the selected file, hover your cursor over the field. The folder location is listed in the tooltip.

If this Scheduler job is stored in a file group Utilities folder, then you can optionally navigate to the report through the **Current File Group** node at the top of the file groups list. When you do this, the path to the file is stored relative to the current file group, which means that it will automatically update when the file group is cloned. This is the recommended method of referencing the report to process when both the Scheduler job and the report belong to the file group.



| Item | Description |
|---|---|
| Override file to process | <p>Optional. Specifies a Scheduler job variable to override the file to process. The override file will then be used for processing instead of the original file. This feature allows you to pass in an alternate file to process, when using Run Event or Raise Event to trigger the Scheduler job for processing.</p> <p>To use a job variable, enter the variable name in curly brackets, such as {MyFile}. When the job is executed, this variable must resolve to a valid file path in the Axiom file repository. Note that it is not valid to leave the variable value blank (the task will <i>not</i> use the original file to process).</p> <p>NOTES:</p> <ul style="list-style-type: none"> • The override feature is only exposed to product developers. It is only visible in client systems if the job is delivered as part of a product package and an override variable is specified in that job. • The File to Process field must point to a valid file for file processing when the override feature is used, even though that file will never actually be processed by the task. If the file to process is missing or invalid, then the task validation will fail. |
| Process Multipass | <p>Specifies whether the report will be run using multipass processing.</p> <ul style="list-style-type: none"> • If this option is selected, multipass processing is performed. This is equivalent to selecting File Output > File Processing > Process File Multipass. • Otherwise, multipass processing is not performed and multipass settings do not display in the task. This is equivalent to selecting File Output > File Processing > Process File. <p>NOTE: If you select Process Multipass, but the file does not have any defined multipass settings, then you must override the blank multipass settings for the file and define them in the equivalent of "advanced mode." If you want to use "basic mode" settings (specify only a source column and Axiom automatically completes the rest of the settings for you), then you should edit the file to define the basic mode multipass settings so that they can be inherited by the task.</p> |
| Enable iterative calculation while processing | <p>Specifies whether iterative calculations are enabled for the file during processing. In most cases you will leave this option disabled.</p> <p>If this option is selected, then iterative calculations are enabled for the file during the Axiom query refresh process. The iterative calculation settings are fixed at a maximum of 100 iterations and a maximum change value of .001.</p> <p>For more information on iterative calculations, see the Microsoft Excel Help.</p> |

► Advanced options

This section only displays if multipass processing is enabled for the task, and the task uses settings that are eligible for parallel processing. Click on the down arrow next to the title to expand the section and view the options.

Parallel processing for file processing tasks is performed based on multipass passes. With certain task settings, multiple passes can be separated into sub-jobs, which can then be processed at the same time (in parallel). This can improve the performance of the task.

For example, imagine that you are multipass processing a file by department. If the task is processed sequentially, then the task would process Dept 100 and finish it, then move to Dept 110 and finish it, etc. When parallel processing is used instead, Depts 100-199 can be separated into one sub-job, Depts 200-299 into another sub-job, etc. Because the sub-jobs are processed in parallel, multiple departments are processed at the same time, so the overall task can complete more quickly.

| Item | Description |
|-----------------------|---|
| Maximum Parallel Jobs | <p>The maximum number of subordinate jobs to run in parallel. The default number is 4.</p> <p>This is the total number of sub-jobs that can be run at the same time for this task. Ultimately the number of sub-jobs that are run in parallel depends on the number of Scheduler threads that have been configured for use at your organization, and the number of Scheduler threads that are currently available (threads that are not processing other higher-priority jobs).</p> |
| Processing Batch Size | <p>The number of multipass passes to include in each sub-job at a time. The default number is 10 for eligible snapshot and export processes, and 7000 for eligible save-to-database processes (save once at end).</p> <p>Passes are determined based on the multipass list of items. For example, if you are processing by department (DEPT.DEPT), then each department is a separate pass. If the batch size is set to 10, then each sub-job would process 10 departments at a time.</p> <p>In most cases, the default settings are sufficient. If you are experiencing lengthy processing times and want to optimize performance, you can adjust this setting as follows: divide the number of passes by the number of available Scheduler threads. For example, if there will be 100 passes and there are 4 Scheduler threads, set the batch size to 25.</p> |

NOTES:

- For save processes, only "save once at end" processes are eligible for parallel processing. In this case, the records to be saved to the database are extracted after each pass to a central temporary table. Once all passes are complete, then all records are saved to the database from the temporary table. Save processes where data is saved directly after each pass are not eligible, because these processes may depend on sequential processing.
- There is no way to disable parallel processing if the task is eligible; however, you can adjust the parallel processing settings if desired.

► File processing settings

Once you have selected a file to process, the file processing settings from that file display within the task as read-only. You can leave the settings as they are, or you can override any setting.

- To override a setting, select the **Override** check box to the right of the setting. The setting becomes editable, and you can change it. The change only applies to the file processing task—the setting remains unchanged within the file.
- If you override a setting, make sure that any related settings make sense in the context of the change. For example, if **File Generation** is set to **Multiple Output Files**, and you override it to be **Single Output File**, then you should also check the **Sheet Names** setting to make sure that you will end up with unique sheet names within the file.

NOTE: If the target file for the task uses **File Collect** or **Batch** processing, then it is not possible to override the settings on the File Collect Configuration Sheet or the Batch Control Sheet.

For more details on file processing settings, see the *Axiom File Setup Guide*.

Note the following requirements when running file processing using Scheduler:

- The **Output Folder** location must be accessible by the Scheduler service user account (for on-premise systems) or the Axiom Cloud Integration Service (for cloud systems). If you specify a network folder location using the Browse button, the location is automatically entered as a UNC path. If you specify a C: drive location, that location will be evaluated as the C: drive of the Scheduler server.
- If the file processing type is **Print**, the Scheduler server(s) must be configured to access the specified printer. This may require the assistance of your IT department.

Job variables can be used in any file processing setting that accepts a typed user input.

► Batch variables

If the file has defined batch variables, you can specify variable values to be used for the file processing task. When the task is executed, any specified variable values are temporarily placed within the file, within the designated cell for that variable value. The file can be set up to use the variable value in some

way during processing.

| Item | Description |
|-----------------|---|
| Variable Names | The names of the variables, as defined in the File Processing Control Sheet for the source file. If no names are listed, then no variables are defined in the file. |
| Variable Values | The variable values to be placed within the file when the file processing task is executed. |

Job variables can be used in the batch variable settings. For example, a job variable can be used as the value for a batch variable.

Import ETL Package task

This task imports data into Axiom (same as executing an import from the **Imports** menu).

NOTE: If the import package is configured to **Ignore lookup and key errors**, then if errors are found the execution status of the job will be Partial Success. This will trigger an email notification if the job is configured to notify only on error.

| Item | Description |
|---------------------------|--|
| Select ETL Import Package | The import package to process. You can select any import that is defined in the current system. |
| Source Filename | <p>The path and name of the source file. This option only applies in the following situations:</p> <ul style="list-style-type: none">• The import is configured to pull data from a source file (instead of a database table).• The import is configured to prompt the user for the source file during execution. <p>If the import is configured to always use the same source file, then that file displays for reference in the Source Filename box, but it is grayed out and cannot be changed.</p> <p>Job variables can be used in this setting.</p> |
| Package Variables | <p>Specifies values for any variables used in the import package.</p> <p>Variables are listed in the right-hand side of the grid. Use the drop-down list next to the variable name to select from the defined set of choices, or type in a value.</p> <p>Job variables can be used in this setting.</p> |

Process Plan Files task

This task processes plan files in a file group. It performs the same actions as the **Process Plan Files** utility available from the file group menu.

The Process Plan Files task uses several tabs to define different options. The available tabs and the options on those tabs depend on the selected **Processing Mode** on the **Options** tab.

- **Options:** Defines the overall processing mode and processing options
- **Plan Files:** Specifies the plan files to process
- **Axiom Queries:** Specifies which Axiom queries to run in plan files (only applies to Normal Processing)
- **Utilities:** Specifies which data source to use for utility processing (only applies to Process with Utilities)
- **Processing Variables:** Defines variables to pass into plan files from Scheduler, and to Scheduler from plan files

► Options tab

The following options are available on the Options tab:

| Item | Description |
|-----------------|---|
| Processing Mode | <p>Select the type of processing to perform:</p> <ul style="list-style-type: none">• Normal Processing: Plan files are opened, refreshed, and saved. You can configure which actions occur.• Process with Utilities: A list of utilities is iteratively processed per plan file. Utilities are opened, refreshed with data for each plan code, and saved. This is primarily intended for processing form-enabled plan files with embedded forms.• Update Persistent Plan Files: Update existing plan files for text, formatting, or formula fixes. This is an advanced feature.• Process with Custom Utility: Plan files are processed using a custom utility provided by Axiom Support. This is an advanced feature. <p>The default processing mode is Normal Processing. However, if the file group has been configured so that utility processing is the default processing mode for that file group, then Process with Utilities is selected by default.</p> |

| Item | Description |
|--|--|
| Select File Group | <p>The file group for which plan files will be processed. You can select any file group or file group alias, including file group scenarios (click Show Scenarios in the Choose File Group dialog to show scenarios in the file group list).</p> <p>If the Scheduler job is stored in a file group Utilities folder, then you can select Use Current File Group to automatically use the file group that the Scheduler job belongs to. This is the recommended approach when the Scheduler job belongs to a file group, so that it will automatically update to point to the current file group when the file group is cloned.</p> <p>NOTE: If the task uses an alias, then you cannot select individual plan files on the Plan Files tab. Only the Use Filter and All options are available.</p> |
| Advanced Options: Worker Batch Size | <p>Optional. Specifies the number of plan files to be processed in each batch. The batch size must be a number between 10 and 100.</p> <p>By default this is left blank, which means that the batch size is automatically calculated based on the number of plan files to be processed divided by the total number of threads on all enabled Scheduler servers. Generally speaking, you should not customize this setting unless you are advised to by Axiom Support.</p> <p>NOTE: Each batch of plan files is processed by a subordinate job. These subordinate jobs are automatically created for the Process Plan Files task and are processed in parallel, dependent on the number of Scheduler threads that are available at any one time.</p> |

Options for Normal Processing mode

If Normal Processing is the selected processing mode, the following additional options are available on the Options tab:

| Option | Description |
|---|---|
| Save document after processing | <p>Specifies whether plan files are saved during processing. This option is selected by default.</p> <p>This option does <i>not</i> cause a save-to-database to be performed—that option must be selected separately.</p> <p>NOTES:</p> <ul style="list-style-type: none"> • If this option is not selected, then the utility will open the file as read-only and will not attempt to acquire the document lock before processing. • If the file group uses virtual plan files, this option does not apply because the plan files cannot be saved. However, if the option is enabled, Axiom will attempt to acquire the document lock before processing, which is not necessary. This option should not be enabled when processing virtual plan files. |
| Run Save To Database on plan files after processing | <p>Specifies whether a save-to-database is performed in plan files during processing. This option is selected by default.</p> <p>This option does <i>not</i> cause the file itself to be saved—that option must be selected separately. It is not required to save the file in order to perform a save-to-database.</p> |
| Create a plan file restore point before processing | <p>If selected, then a plan file restore point will be created before processing begins. This option is not selected by default.</p> <p>Restore points can be used to restore plan files to the state they were in before changes were made.</p> <p>NOTE: If the file group uses virtual plan files, this option does not apply. Plan files are not saved and therefore restore points are irrelevant.</p> |

Options for Process with Utilities

If **Process with Utilities** is the selected processing mode, there are no additional options on the Options tab.

Plan files are not saved when using Process with Utilities, and plan file restore points are not created. When using this mode, the processing is being performed in the utility files, not in the plan files, so it is not necessary to save the plan files. Additionally, in most cases the plan files used with this mode are virtual form-enabled plan files, so the save and restore options are irrelevant.

Options for Update Persistent Plan Files

If **Update Persistent Plan Files** is the selected processing mode, the following additional option is available on the Options tab:

| Option | Description |
|-------------|--|
| Report File | Click the Browse button to select the report file that is configured with the PlanFileReconfig_ControlSheet. This file must be saved in the Reports Library. This control sheet contains the settings that will be applied to plan files during processing. |

Plan files are always saved when using this processing option, and plan file restore points are always created before processing. A save-to-database is not performed in this mode, so if you need to save data, you should process plan files using Normal Processing after you have verified the results of the plan file update.

Options for Process with Custom Utility

If **Process with Custom Utility** is the selected processing mode, the following additional options are available on the Options tab:

| Item | Description |
|--------------|--|
| Report File | Click the Browse button to select the Microsoft Excel spreadsheet file that contains the VBA custom utility. The file must be saved in the Reports Library. |
| VBA Module | Select the VBA module to run as part of this utility. The drop-down list shows the VBA modules available in the selected file. |
| VBA Function | Select the VBA function to run as part of this utility. The drop-down list shows the VBA functions available in the selected module. |

Plan files are always saved when using this processing option, and plan file restore points are always created before processing. A save-to-database is not performed in this mode, so if you need to save data, you should process plan files using Normal Processing after you have verified the results of the custom utility processing.

► Plan Files tab

On the **Plan Files** tab, specify the plan files that you want to process. There are three different options that you can use to specify the plan files: **Choose from list**, **Use filter**, and **All**. You should use the option that corresponds to how many plan files you want to process—all plan files, or a subset of plan files. If you want to process a subset of plan files, you can select individual files to process or you can use a filter to define the subset.

NOTES:

- If a plan file is locked by another user when the task is executed, then processing for that file will fail. Failures are noted in the result history for the job.
- If a plan file has not yet been created for a particular plan code, then that plan code will not display in this list and will be ignored when processing. Scheduler does not support creating plan files as part of the Process Plan Files task (you must use the separate Create Plan Files task for this purpose).
- If the file group uses a **Show on List** column, then any plan code that is set to **False** will not display in the plan file list and will be ignored when processing.

Process all plan files

To process all plan files, select **All**. The list of all plan files is generated each time the Scheduler task is executed, so that if new plan files have been added then those new plan files will be included in the processing (the reverse is also true if any plan files have been removed).

Alternatively, you can select **Choose from list** and then select the check box in the column header, causing all current plan codes to be selected. However, in this case the list of selected plan codes is fixed and therefore will not automatically adjust for any future changes.

Process selected plan files

To process certain plan files, select **Choose from list**, and then select the check boxes for the plan files that you want to process. When the Scheduler task is executed, Axiom will process only the selected plan files.

To find the plan files you are looking for, you can sort, filter, and group the list using standard Axiom grid features. You can show additional columns and hide columns by right-clicking in the column header. If you have filtered the list, you can select the check box in the header to select only the plan files that currently display in the dialog.

NOTE: This option is not available if the file group for the task is an alias. This is because the list of plan files could change when the alias target changes.

Process a filtered set of plan files

To use a filter to process a subset of plan files, select **Use Filter**. When the Scheduler task is executed, Axiom will process only the plan files that meet the filter.

You can use the Filter Wizard to create the filter, or you can manually type a filter criteria statement into the filter box. The filter must use the plan code table or a lookup table. For example:

`DEPT.Region='US West' where Dept is the plan code table.`

Once you have entered a filter, you can click **Refresh plan file list** to show the plan files that currently match the filter. The refresh feature is intended to help you determine whether you have defined the filter correctly.

You can also use a job variable for the filter. For example, you can define a job variable named "filter" and then place the text {filter} in the filter box. This is intended for use when running Scheduler jobs via RunEvent. If a variable value is specified when the event is triggered, such as the value `dept.region='west'`, then that filter statement will replace the {filter} variable and will be used to determine the list of plan files to be processed.

NOTE: If you use a variable, and you leave the default value for that variable blank within the **Job Variables** tab, then all plan codes will be processed if no value is passed by the RunEvent function. You may want to define a default filter that results in no values (such as 1=0), so that plan files are only processed if a valid filter value is passed.

► Axiom Queries

On the **Axiom Queries** tab, select the queries that you want to run in the plan files. By default, all listed queries are selected. This tab only applies when using **Normal Processing** mode.

If you do not want to run a particular query, you can clear the check box. You can select or clear individual check boxes, or you can use the check box in the header to select or clear all queries currently displayed in the list. You can sort, filter, and group the list using standard Axiom grid functionality.

| Options Plan Files Axiom Queries | | | | | |
|---|----------|-----------|---|-----------------|---------|
| Active Axiom Queries for selected Plan Files are shown in the list below. Selected Axiom Queries will be run when the related Plan Files are processed. | | | | | |
| <input checked="" type="checkbox"/> | Template | Worksheet | Axiom Query | Refresh On Open | Dynamic |
| <input checked="" type="checkbox"/> | Master | Stat_Rev | AQ1: Stat_Rev | False | False |
| <input checked="" type="checkbox"/> | Master | Stat_Rev | AQ2: NetRevSection | False | True |
| <input checked="" type="checkbox"/> | Master | Stat_Rev | AQ3: Forecast | False | True |
| <input checked="" type="checkbox"/> | Master | Stat_Rev | AQ4: ColHide On Open | True | False |
| <input checked="" type="checkbox"/> | Master | Stat_Rev | AQ5: Statistics On Open | True | False |
| <input checked="" type="checkbox"/> | Master | JobCode | AQ1: Labor Configuration Driver On Open | True | False |
| <input checked="" type="checkbox"/> | Master | JobCode | AQ2: Labor Configuration Driver On Open | True | False |

Example Axiom Queries tab

The list of Axiom queries is based on the source templates that were used to create the plan files. Only Axiom queries that meet the following criteria are eligible for selection:

- **Active** is set to **On**, or the setting uses a formula.
- **Refresh during document processing** is set to **On**.

If a query uses a formula for the Active setting, this means the query is dynamic and may or may not be run, depending on how the formula resolves in each plan file to be processed. When a particular plan file is processed, each selected query will be evaluated based on the current settings in that plan file. If both **Active** and **Refresh during document processing** are **On** for that plan file, then the query will be run. If either or both settings are **Off** for that plan file, the query will not be run. You can tell whether a query is dynamic or not by looking at the **Dynamic** column in the query list.

If a query is *not* selected on this tab, then that query will not be run in any plan files during processing, regardless of whether **Active** or **Refresh during document processing** are enabled in the plan file.

The plan file selection on the **Plan Files** tab affects the Axiom query list as follows:

- If you have selected individual plan files, then only the eligible queries for the source templates of the selected plan files are shown.
- If you have selected **All** or **Use Filter**, then all eligible queries for all used templates are shown. If the file group has templates that have not been used to create any plan files, then those templates are not included in the list.

The listed queries are identified by template, worksheet, and query name. The following additional properties are also listed for each query:

- **Refresh On Open**: Indicates whether the Axiom query is configured to refresh automatically when the file is opened. This is for information purposes only, to help you determine whether the query needs to be included in the processing. The Refresh on Open status is ignored by Process Plan Files—if the query is selected it will be run along with the other selected queries, and if it is not selected it will not be run.
- **Dynamic**: Indicates whether the query is dynamically enabled. True means that the query uses a formula for the **Active** setting.

NOTE: If a query is listed on this tab but it is grayed out and unavailable for selection, that means that although the query is active (either directly or dynamically), the query is not eligible to be run using Process Plan Files (because the setting **Refresh during document processing** is set to **Off**). This query is listed for your information only, so that you understand the query cannot be run as part of the process.

► Utilities tab

On the **Utilities** tab, select the ProcessPlanFileUtilities data source to use during processing. This data source determines which utility files are processed and the processing order. This tab only applies when using **Process with Utilities** mode.

The list below shows the selected utilities data source for each template in the currently selected plan files.

| Template Name | Utilities Data Source |
|---------------|-----------------------|
| FormTemplate | Set1 |

Example Utilities tab

For each template listed, use the **Utilities Data Source** field to select the data source to use for plan files created from that template.

- If the template only has one data source, that data source is selected.
- If the template has multiple data sources, then the data source marked as the default data source is selected by default. If desired, you can use the drop-down list to select a different data source.

When plan files are processed, Axiom reads the specified data source in each plan file to determine the utilities to be processed for that plan file.

The plan file selection on the **Plan Files** tab affects the Utilities list as follows:

- If you have selected individual plan files, then only the templates used to create the selected plan files are shown.
- If you have selected **All** or **Use Filter**, then all used templates are shown. If the file group has templates that have not been used to create any plan files, then those templates are not included in the list.

► Processing Variables

This tab can be used to define variables to pass into plan files before processing begins, and to pass variables back to the Scheduler job after processing has been performed. This tab is optional and is only used in special situations.

Pre-Processing Document Variables

This section can be used to pass document variables into plan files before processing. This can impact the processing of plan files if the files are configured to use the variable values in some way.

For each pre-processing document variable, you can specify a variable name and a variable value. The plan files must be set up with GetDocumentInfo functions that return the values for the specified variables.

- + To add a variable, click the Add button to add a row to the list. Complete the settings for the variable as described below.
- ✕ To remove a variable, select the variable in the list and then click the Remove button. Only one variable can be selected at a time.

To edit the variable settings, double-click the applicable cell to make the cell contents editable. When you are finished editing, you can press the Enter key or Tab key to exit the cell, or click outside of the cell.

| Item | Description |
|----------------|--|
| Variable Name | The name of the variable. Do not enclose the variable name in curly brackets (you are not <i>using</i> the variable here, you are defining its value). |
| Variable Value | <p>The value of the variable. The value can be a "hard-coded" value, or it can be a job variable that will be resolved at time of processing.</p> <p>If you use a job variable to define the value, the job variable must be enclosed in curly brackets.</p> |

Pre-Processing Workbook Variables

This section can be used to pass values into plan files before processing. This can impact the processing of plan files if the files are configured to use the values in some way.

For each pre-processing variable, you can specify a workbook location to place the value, and the value to be placed.

| Item | Description |
|-------------------|---|
| Workbook Location | <p>The location in the workbook for the value to be placed. Any existing value in this location will be overwritten for the duration of the processing. If the file is saved as part of the processing, then the value will be saved in the file.</p> <p>The location can be specified using <i>SheetName!CellRef</i> syntax (for example: <i>Report!A13</i>), or by using a named location in the file.</p> |
| Formula | <p>The value to be placed in the specified workbook location. The value can be a "hard-coded" value, or a formula, or a job variable that will be resolved at time of processing.</p> <p>If the value is a formula, the formula is placed into the target cell and calculated in the plan file. The formula can be any formula that would be valid within a spreadsheet in the Axiom client. This includes using Excel functions and Axiom functions. The formula can also use job variables, which will be resolved before placing the formula in the target cell.</p> |

The specified location and value will apply to all plan files being processed by the task. If you are going to use pre-processing variables, the location should be predefined in the template and therefore available to all plan files built using that template. If the plan files will be built using multiple templates, then all templates should be set up with the same designated location, or you should set up separate processing tasks based on template type.

Post-Processing Workbook Variables

This section can be used to pass a value from plan files back to the Scheduler job after processing has been performed. This can impact the processing of subsequent tasks in the job if those tasks are configured to use the value in some way.

For each post-processing variable, you can specify the location in the workbook to find the value, and the job variable to use that value.

NOTE: If this task processes multiple plan files, the resulting variable value will be from the last file that was processed.

| Item | Description |
|-------------------|--|
| Workbook Location | <p>The location in the workbook to find the value to be passed to Scheduler. This value will become the value for the assigned job variable for the duration of executing the current job (unless a later process within the same job overwrites the value for the same job variable).</p> <p>The location can be specified using <i>SheetName!CellRef</i> syntax (for example: <i>Report!A13</i>), or by using a named location in the file.</p> |
| Job Variable | <p>The job variable that you want to use the value in the specified workbook location. Do not enclose the variable name in curly brackets (you are not <i>using</i> the variable here, you are simply referencing the variable name).</p> <p>If the job variable does not already exist in the job (on the Job Variables tab), then it will be created. However, in most cases you will want the variable to be already set up with a default value, so that the job does not have validation errors that prevent saving.</p> |

The specified location and job variable will apply to all plan files being processed by the task. If you are going to use post-processing variables, the location should be predefined in the template and therefore available to all plan files built using that template. If the plan files will be built using multiple templates, then all templates should be set up with the same designated location, or you should set up separate processing tasks based on template type.

Even though the task may process many plan files, only the job variable value from the last-processed plan file will be used. The plan files must be set up so that all plan files result in the same value after processing, or else your results will vary depending on which plan file was the last file to be processed.

Process Document List task

This task processes a user-defined set of documents. The process operation always calculates the files. In addition, you can opt to run Axiom queries in the files, process alerts in the files, and then perform a save-to-database and/or save the files.




You can process any Axiom-managed Excel files by using this task. The primary intent of the task is to process files such as driver files or report utilities. For example, you may be using Axiom queries and GetData functions in your driver files that need to be updated regularly. Rather than opening, refreshing, and saving each driver file, you can use this task to define the set of files and schedule processing.

NOTES:

- Generally speaking, plan files should not be processed using this task. Instead, the [Process Plan Files](#) task should be used.
- This task does not perform *file processing* actions on the file. File processing can be set up for report files and driver files, and can be used to perform actions such as file delivery, using standard or multipass processing. If you want to perform file processing using Scheduler, use the [File Processing](#) task.

► Documents to process

Specify the documents to be processed when the task is run. Documents are processed sequentially in the order listed.


-  To add a document, click the Add button. In the **Axiom Explorer** dialog, select the file or files that you want to add, and then click **Open**.
-  To remove a document, select the document in the list and then click the Remove button. Only one document can be selected at a time.
-  To change the order of documents, select the file in the list and then click the arrow buttons to move the file up or down.

Only Axiom-managed Excel files are valid to be processed in the task.

Selecting a document using a file group alias


You may want to specify the document to process using a file group alias, so that the task does not have to be manually updated after rolling over to a new file group for a new year of planning. When you are selecting a document in the Axiom Explorer dialog, you can expand the file group alias to see all of the files in the current target of the alias.

For example, in the following screenshot, the file group alias Current Budget is expanded. Because the alias currently points to the file group Budget 2018, the folders and files under the alias are from Budget 2018. When you select a file or folder, you can see the real file path in the top of Axiom Explorer.



Selecting a document to process using a file group alias

When you select a document underneath an alias, the file path is written using alias syntax, so that the Scheduler task will look for the document within the current target of the alias. This path is visible in the tooltip that displays when you hover your cursor over a selected document.



File path using file group alias syntax

When the file group alias is updated to point to a new file group, the Scheduler task will use the file in the new file group automatically. If the file does not exist in the new file group, the task will fail with an error.

► Processing Options

By default, both options are selected. Axiom queries are refreshed before the save-to-database occurs.

If neither option is selected, then the files are calculated and then saved.

| Item | Description |
|---|---|
| Perform all enabled Axiom Queries in selected workbooks | <p>If this option is selected, then all eligible Axiom queries in all selected files will be refreshed when the task is run. This option is selected by default.</p> <p>Axiom queries are eligible for processing if they are active and use either of the following refresh behaviors: Refresh on File Open and Refresh During Document Processing.</p> |
| Enable iterative calculation while processing | <p>Specifies whether iterative calculations are enabled for the file during processing. In most cases you will leave this option disabled.</p> <p>If this option is selected, then iterative calculations are enabled for the file during the Axiom query refresh process. The iterative calculation settings are fixed at a maximum of 100 iterations and a maximum change value of .001.</p> <p>For more information on iterative calculations, see the Microsoft Excel Help.</p> |
| Save document after processing | <p>If selected, then files will be saved after processing. This option is selected by default. The user executing the task must have Read/Write access to the files.</p> <p>This option does <i>not</i> cause a save-to-database to be performed—that option must be selected separately.</p> <p>NOTE: If this option is not selected, then the utility will open the file as read-only and will not attempt to acquire the document lock before processing.</p> |
| Run Save To Database on plan files after processing | <p>If selected, then a save-to-database will be performed after processing. This option is selected by default. The user executing the task must have the Allow Save Data permission to the files.</p> <p>This option does <i>not</i> cause the file itself to be saved—that option must be selected separately. It is not required to save the file in order to perform a save-to-database.</p> |
| Process alerts in selected workbooks | <p>If selected, then alerts in the file will be processed. The file must contain an Alert Control Sheet and one or more alerts must be defined in the file.</p> <p>If Axiom queries are enabled for processing as well, the queries will be run before alerts are processed.</p> |

► Pre-Processing Document Variables

This section can be used to pass document variables into the target files before processing. This can impact processing if the files are configured to use the variable values in some way, such as to filter an Axiom query.


For each pre-processing document variable, you can specify a variable name and a variable value. The target file must be set up with GetDocumentInfo functions that return the values for the specified variables.

- ➕ To add a variable, click the Add button to add a row to the list. Complete the settings for the variable as described below.
- ✕ To remove a variable, select the variable in the list and then click the Remove button. Only one variable can be selected at a time.

To edit the variable settings, double-click the applicable cell to make the cell contents editable. When you are finished editing, you can press the Enter key or Tab key to exit the cell, or click outside of the cell.

| Item | Description |
|----------------|---|
| Variable Name | The name of the variable. Do not enclose the variable name in curly brackets (you are not <i>using</i> the variable here, you are defining its value). |
| Variable Value | The value of the variable. The value can be a "hard-coded" value, or it can be a job variable that will be resolved at time of processing. If you use a job variable to define the value, the job variable must be enclosed in curly brackets. |




The following example screenshot defines the value for the document variable Dept. The value is defined using the value of a Scheduler job variable {Dept}. A value for Dept can be passed into the Scheduler job when the job is started, and then passed into the target file using the document variable.



► Pre-Processing Workbook Variables

This section can be used to pass values into the file before processing. This can impact processing if the files are configured to use the values in some way. For each pre-processing variable, you can specify a workbook location to place the value, and the value to be placed.

NOTE: The specified value will be placed in all files listed to process. Therefore, unless all files are set up to use the same location and the same way of deriving the value, you will need to create a separate Process Document List task for each file. If you are not using pre-processing variables, then this does not matter.

-  To add a variable, click the Add button to add a row to the list. Complete the settings for the variable as described below.
-  To remove a variable, select the variable in the list and then click the Remove button. Only one variable can be selected at a time.
-  To change the order of variables, select the variable in the list and then click the arrow buttons to move the variable up or down.

To edit the variable settings, double-click the applicable cell to make the cell contents editable. When you are finished editing, you can press the Enter key or Tab key to exit the cell, or click outside of the cell.

| Item | Description |
|-------------------|---|
| Workbook Location | <p>The location in the workbook for the value to be placed. Any existing value in this location will be overwritten for the duration of the processing. If the file is saved as part of the processing, the value will be saved in the file.</p> <p>The location can be specified using <i>SheetName!CellRef</i> syntax (for example: <i>Report!A13</i>), or by using a named location in the file.</p> |
| Formula | <p>The value to be placed in the specified workbook location. The value can be a "hard-coded" value, or a formula, or a job variable that will be resolved at time of processing.</p> <p>If the value is a formula, the formula is placed into the target cell and calculated in the target file. The formula can be any formula that would be valid within a spreadsheet in the Axiom client. This includes using Excel functions and Axiom functions. The formula can also use job variables, which will be resolved before placing the formula in the target cell.</p> |

► Post-Processing Workbook Variables

This section can be used to pass a value from the file back to Scheduler after processing has been performed. This can impact the processing of subsequent tasks in the job if those tasks are configured to use the value in some way. For each post-processing variable, you can specify the location in the workbook to find the value, and the job variable to use that value.

NOTES:

- If this task processes multiple documents, the resulting variable value will be from the last document processed.
- If the task is run as a sub-job, then the post-processing variable is not passed back to the "parent" job. The task must be executed as a regular task within the job in order to pass the variable value back to the job.

- + To add a variable, click the Add button to add a row to the list. Complete the settings for the variable as described below.
- ✕ To remove a variable, select the variable in the list and then click the Remove button. Only one variable can be selected at a time.
- ▲ To change the order of variables, select the variable in the list and then click the arrow buttons to move the variable up or down.

To edit the variable settings, double-click the applicable cell to make the cell contents editable. When you are finished editing, you can press the Enter key or Tab key to exit the cell, or click outside of the cell.

| Item | Description |
|-------------------|---|
| Workbook Location | <p>The location in the workbook to find the value to be passed to Scheduler. This value will become the value for the assigned job variable for the duration of executing the current job (unless a later process within the same job overwrites the value for the same job variable).</p> <p>The location can be specified using <i>SheetName!CellRef</i> syntax (for example: <i>Report!A13</i>), or by using a named location in the file.</p> |
| Job Variable | <p>The job variable that you want to use the value in the specified workbook location. Do not enclose the variable name in curly brackets (you are not <i>using</i> the variable here, you are simply referencing the variable name).</p> <p>If the job variable does not already exist in the job (on the Job Variables tab), then it will be created when the job is executed. However, in most cases you will want the variable to be already set up with a default value, so that the job does not have validation errors that prevent saving.</p> |

Process Template List task

This task processes a user-defined list of file group templates. During processing, any Axiom queries with **Refresh during template processing** enabled are executed and time-stamped, and then the template files are saved.

The primary purpose of this task is to enable use of *time-stamped Axiom queries* with virtual plan files. Because virtual plan files are re-created from template each time they are accessed, Axiom queries cannot be time-stamped within the plan files. Virtual plan files can use the time stamp from the template, but under normal circumstances, Axiom queries are not time-stamped when they are run in templates. However, when Axiom queries are run during template processing, the **Last refresh time** for the query is updated, which means that the queries can be configured to only run if the primary table has changed.

To use this task to enable time-stamped Axiom queries for virtual plan files, do the following:

- In the template, enable **Refresh only if primary table changed since last refresh** and **Refresh during template processing** for the Axiom queries that you want to be time-stamped.
- In Scheduler, create a job with a **Process Template List** task and add the template to the task. Define a scheduling rule for the job as appropriate. For example, you might want the template to be processed nightly.

When the template is processed, the designated Axiom queries will be run if the primary table has changed, and the time stamps are updated. When a virtual plan file that uses this template is opened, the queries will not be run again if the primary table has not changed.

This task should only be used to process Axiom queries that meet the requirements of time-stamped queries.

► Templates to process

Specify the templates to be processed when the task is run. Templates are processed sequentially in the order listed. If you have multiple templates to process (in the same or different file groups), you can run them all in the same task.

- + To add a template, click the Add button. In the **Axiom Explorer** dialog, select the file or files that you want to add, and then click **Open**.
- ✕ To remove a template, select the file in the list and then click the Remove button. Only one file can be selected at a time.
- ▲ ▼ To change the order of templates, select the file in the list and then click the arrow buttons to move the file up or down.


Normal template behavior rules apply during processing. For example, save-to-database and action codes are not run in templates. The only exception to normal template behavior during this task is that any executed Axiom queries will be time stamped.

Selecting a template using a file group alias

You may want to specify the template to process using a file group alias, so that the task does not have to be manually updated after rolling over to a new file group for a new year of planning. When you are selecting a template in the Axiom Explorer dialog, you can expand the file group alias node to see all of


the files in the current target of the alias.

For example, in the following screenshot, the file group alias Current Budget is expanded. Because the alias currently points to the file group Budget 2018, the folders and files under the alias are from Budget 2018. When you select a file or folder, you can see the real file path in the top of Axiom Explorer.



Selecting a template to process using a file group alias

When you select a template underneath an alias, the file path is written using alias syntax, so that the Scheduler task will look for the template within the current target of the alias. This path is visible in the tooltip that displays when you hover your cursor over a selected template.



File path using file group alias syntax

When the file group alias is updated to point to a new file group, the Scheduler task will use the file in the new file group automatically. If the file does not exist in the new file group, the task will fail with an error.

Purge System Data task

The Purge System Data task is intended to clean up old data in your system, to help keep your system running efficiently.

NOTE: Scheduler automatically creates a system job for this task (System.SystemDataPurge), which administrators can edit as needed.

This task purges the following data when it is run:

- Scheduler job result history
- Scheduler and system email notifications
- System temp table data
- Audit history
- Alerts

For each category of data, you can specify a number of days of data to keep when the task is run. All results older than the specified number of days will be deleted. Note that 0 days means that no data is purged for that category.

| Section | Item | Description |
|-------------------|---|---|
| Scheduler Results | Number of days to keep result history | The number of days of job result history to keep when the task is run. By default, this is set to 15 days. |
| SMTP Messages | Number of days to keep delivered messages and attachment data | The number of days of delivered message data to keep when the task is run. By default, this is set to 15 days. |
| Temporary Tables | Number of days to keep temp table data | The number of days of temp table data to keep when the task is run. By default, this is set to 15 days. |
| Audit History | Number of days to keep system history | <p>The number of days of system audit history to keep when the task is run. By default, this is set to 15 days.</p> <p>"System history" encompasses all audit data—including prior document versions and deleted documents—except table audit data.</p> |
| Table History | Number of days to keep table history | <p>The number of days of table audit history to keep when the task is run. By default, this is set to 15 days.</p> <p>Table audit data is tracked for tables where Audited is set to True.</p> |

| Section | Item | Description |
|---------|-------------------------------|--|
| Alerts | Number of days to keep alerts | The number of days of alerts to keep when the task is run. By default, this is set to 60 days. |

Job variables can be used in all of these settings.

Each purge routine in the task is limited to purging a specific number of rows at a time (50000). If the number of rows to be purged exceeds this limit, then the excess data is retained until the next time the task is run. If you notice data in the database that you expected to be purged, most likely the amount of data to be purged exceeded the limit, and the data will be purged next time the task is run.

► Other purged data

This task also cleans up the following items in your system:

- Deleted columns. When a column is deleted from a table in Axiom, the column is immediately deleted from the associated view (which prevents it from being accessed in the system), but it remains in the base table. This task finishes the process of removing obsolete columns from the base tables.
- Orphaned user folders. If the system contains any user folders that do not match up with existing users, these folders are deleted. Although user folders are deleted when a user is deleted from security, orphaned user folders can result from other processes, such as migrating a system between different management databases.

These items are not associated with any specific task settings; the delete process is performed whenever the task is executed.



Raise Event task

The Raise Event task can be used to trigger other Scheduler jobs for execution, using a named event handler. This task has one required setting:

| Item | Description |
|------------|--|
| Event Name | <p>Enter the name of the event that you want to raise for execution. This name must match a defined event handler name in one or more other Scheduler jobs.</p> <p>When this task is run, it looks for any jobs that contain the specified event handler name. These jobs are added to the schedule and are eligible to be processed immediately, depending on Scheduler thread availability and any other higher-priority jobs already in the queue.</p> <p>NOTE: It is not possible to specify a file group context for the event handler when using Raise Event. Axiom will run all jobs that contain the specified event handler name, regardless of whether the event handler is associated with a file group.</p> |

Event Variables

This section can be used to pass variables into the jobs triggered by the event handler. If the jobs are configured to use the variables, these values can impact how the jobs are processed.

-  To add a variable, click the Add button to add a row to the list. Complete the settings for the variable as described below.
-  To remove a variable, select the variable in the list and then click the Remove button. Only one variable can be selected at a time.

For each variable, you can specify a variable name and a variable value. To edit the variable settings, double-click the applicable cell to make the cell contents editable. When you are finished editing, you can press the Enter key or Tab key to exit the cell, or click outside of the cell.

| Item | Description |
|----------------|--|
| Variable Name | The name of the variable. Do not enclose the variable name in curly brackets (you are not <i>using</i> the variable here, you are defining its value). |
| Variable Value | <p>The value of the variable. The value can be a "hard-coded" value, or it can be a job variable that will be resolved at time of processing.</p> <p>If you use a job variable to define the value, the job variable must be enclosed in curly brackets.</p> |

Run Scheduler Job task

This task runs a specified Scheduler job as a subordinate job within the current job. The job containing the Run Scheduler job task is the parent job, and the target job for the task is the child job.

By default, the parent job waits until the child job is complete before continuing to the next task in the parent job. This means that tasks after the Run Scheduler Job task can be reference the results of the child job. For example, the child job may perform a save-to-database. The subsequent tasks in the parent job can access the data saved by the child job.

► Task Control options

When you create the Run Scheduler Job task, the options in the **Task Control** section are pre-set as follows:

- The option **Create a Subordinate Job for this Task** is grayed out. This is because the target job is always run as a subordinate job.
- The option **Wait for all Subordinate Jobs to complete before proceeding to the next Task** is enabled by default. This means that tasks after the Run Scheduler Job task can be dependent on the target job and reference the results of that job. If you disable this option, then the parent job will continue to the next task in the job immediately after creating the subordinate job—it will not wait for the subordinate job to complete.

Keep in mind that if you want the parent job to stop processing tasks if the target job fails, then **Wait for all Subordinate Jobs to complete before proceeding to the next Task** must be enabled for the task and **If this Task fails, continue executing subsequent Tasks** must be disabled for the task. This is the default configuration.

If needed, you can use the option **Process task only if the value of this expression is true** to detect whether a subsequent task in the parent job should be processed. For example, if you know that the child job saves a particular value to the database, you can check for the existence of that value to determine whether to process the task. For more information on using this option, see [Conditionally processing tasks in a job](#).

► Target Scheduler job

In the **Task Details** section, use the Browse button to select the target **Scheduler Job**. You can select any job that you have access to in the Scheduled Jobs Library.

When the Run Scheduler Job task is executed, it creates one or more subordinate jobs as needed to execute the tasks in the target Scheduler job. As long as **Wait for all Subordinate Jobs to complete before proceeding to the next Task** remains enabled in the Task Control options, the parent job waits for all subordinate jobs to be completed before moving on to the next task in the parent job.

NOTE: The user executing the job does not need to have security access to the target Scheduler job for Run Scheduler Job. It is assumed that if the user can execute the parent job, the user should be able to execute the target job.

► Child Job Values

If the target job for the Run Scheduler Job task has defined job variables, those variables and their default values are listed in this section. The default values are determined as follows:

- If the parent job and the child job have a variable with the same name, the default value is the value defined in the parent job. This value will be passed to the child job and used when the child job is run.
- Otherwise, the default value is the value defined in the child job.

To override a variable value, select the **Override** check box and then click inside the **Override Value** field to enter a value. You can enter a hard-coded value or use a job variable from the parent job. Enter the variable name in curly brackets to use that variable's value as the override value.

For example, imagine that both the parent job and the child job have a variable of `{Dept}`. In the parent job, the value of `{Dept}` is set to 20000, and in the child job the value is set to 40000. The Run Scheduler Job task will display the parent value of 20000 as the default value, and that value will be used when the child job is executed.

Now imagine that the parent job has a variable of `{StartDept}` set to 20000, and the child job has a variable of `{Dept}` set to 40000. In this case, the Run Scheduler Job task will display the value of `{Dept}` as defined in the child job (40000). If you want to use the parent job value for `{StartDept}` instead, then you must select the Override check box and enter `{StartDept}` as the Override Value. Now the value of `{Dept}` in the child job will be overridden and set to 20000.

SMTP Message Delivery task

This task delivers email notifications for Scheduler jobs.

NOTE: Scheduler automatically creates a system job for this task (System.SMTPMessageDelivery), which administrators can edit as needed.

| Item | Description |
|--------------------------------|--|
| Server name | The server name of the SMTP email server. |
| Port number | The port number for the SMTP email server. By default, the port is 587, but you can specify a different port number if needed. |
| Server requires authentication | Select this check box if the SMTP email server requires authentication. If selected, type a Username and Password . |

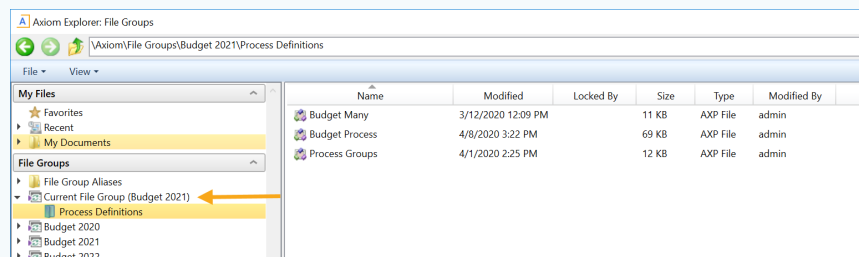
| Item | Description |
|-----------|---|
| Test Mode | <p>Specifies whether the task is run in test mode. If this check box is selected, the task verifies that it can successfully connect to the SMTP server to send email notifications, but no emails are actually sent.</p> <p>For the System.SMTPMessageDelivery job, new systems are automatically set to test mode. If you restore a database, the restore process also sets the system job to test mode. You must disable test mode before any emails will be sent.</p> |

Start Process task

This task starts a process for Process Management. You can use this task to automatically start a process at a specific point in time, including recurring schedules (such as to automatically start a monthly process).

This task can be used to start a generic process definition or a plan file process definition.

| Item | Description |
|------------------|---|
| Process to start | <p>The process definition to start. Click the Browse button to select the process definition file. You can select any process definition that you have access to within the Process Definition Library or within a file group Process Definitions folder.</p> <p>If the Scheduler job is stored in a file group Utilities folder, then you can optionally navigate to the process definition file through the Current File Group node at the top of the file groups list. When you do this, the path to the file is stored relative to the current file group, which means that it will automatically update when the file group is cloned. This is the recommended method of referencing the process definition file when the Scheduler job belongs to a file group.</p> |



| Item | Description |
|--|---|
| Restart process if it is already running | <p>Specifies whether the Scheduler task will restart the process if it is already running, or if the process will be left as is.</p> <ul style="list-style-type: none"> • Select this option if you want to start the target process regardless of whether it is already running. The current process instance will be aborted and a new process instance will start over at step 1. This option is selected by default. • Clear this option if you want to leave the existing process instance running. In this case, the Scheduler task will take no action if the target process is already running. |

Web Report Processing task

This task performs multipass processing on a web report. The report is processed over a dimension with a filter automatically applied to limit the data in the report to the current dimension value. The result of each pass is either a PDF or Excel output file, which can be saved to a designated folder location and/or emailed.

IMPORTANT: Only web reports created from product-delivered templates can be processed by this task. Consult the separate product documentation for more information on any web report templates delivered with your product. Web reports that were created in the Report Builder cannot be processed at this time.

► General task properties

The general task properties at the top of the task settings determine the report to process.

| Item | Description |
|-----------------|--|
| File to Process | <p>The web report to process for the task. Click the Browse button to open the Axiom Explorer dialog, and then select a report to process. You can select any template-based web report that you have access to within the Reports Library.</p> <p>Only one report can be selected for each File Processing task. If you want to process multiple reports, you can add multiple File Processing tasks to the Scheduler job.</p> <p>TIP: Once the file is selected, only the file name displays in the task. If you want to know the folder location of the selected file, hover your cursor over the field. The folder location is listed in the tooltip.</p> |

| Item | Description |
|--------------------------|---|
| Override file to process | <p>Optional. Specifies a Scheduler job variable to override the file to process. The override file will then be used for processing instead of the original file. This feature allows you to pass in an alternate file to process, when using Run Event or Raise Event to trigger the Scheduler job for processing.</p> <p>To use a job variable, enter the variable name in curly brackets, such as {MyFile}. When the job is executed, this variable must resolve to a valid file path in the Axiom file repository. Note that it is not valid to leave the variable value blank (the task will <i>not</i> use the original file to process).</p> <p>NOTES:</p> <ul style="list-style-type: none"> • The override feature is only exposed to product developers. It is only visible in client systems if the job is delivered as part of a product package and an override variable is specified in that job. • The File to Process field must point to a valid file for file processing when the override feature is used, even though that file will never actually be processed by the task. If the file to process is missing or invalid, then the task validation will fail. |

Advanced options

When the task is configured to output multiple files (**File Generation** is set to **Multiple Output Files**), then multiple passes can be separated into sub-jobs, which can then be processed at the same time. This parallel processing can improve the performance of the task.

For example, imagine that you are multipass processing a report by department. If the task is processed sequentially, then the task would process Dept 100 and finish it, then move to Dept 110 and finish it, and so on. When parallel processing is used instead, Depts 100-199 can be separated into one sub-job, Depts 200-299 into another sub-job, etc. Because the sub-jobs are processed in parallel, multiple departments are processed at the same time, so the overall task can complete more quickly.

| Item | Description |
|-----------------------|---|
| Maximum Parallel Jobs | <p>The maximum number of subordinate jobs to run in parallel. The default number is 4.</p> <p>This is the total number of sub-jobs that can be run at the same time for this task. Ultimately the number of sub-jobs that are run in parallel depends on the number of Scheduler threads that have been configured for use at your organization, and the number of Scheduler threads that are currently available (threads that are not processing other higher-priority jobs).</p> |


| Item | Description |
|-----------------------|---|
| Processing Batch Size | <p>The number of multipass passes to include in each sub-job at a time. The default number is 10.</p> <p>Passes are determined based on the multipass list of items. For example, if you are processing by department (DEPT.DEPT), then each department is a separate pass. If the batch size is set to 10, then each sub-job would process 10 departments at a time.</p> <p>In most cases, the default settings are sufficient. If you are experiencing lengthy processing times and want to optimize performance, you can adjust this setting as follows: divide the number of passes by the number of available Scheduler threads. For example, if there will be 100 passes and there are 4 Scheduler threads, set the batch size to 25.</p> |


NOTE: There is no way to disable parallel processing if the task is eligible; however, you can adjust the parallel processing settings if desired.

► Report processing properties

The report processing properties in the middle of the task settings determine the output of the task.

| Item | Description |
|---------------------|---|
| Processing Type | <p>Select one of the following to determine the output format of each pass:</p> <ul style="list-style-type: none"> • Export to Excel (default): The contents of the report are exported to a spreadsheet (XLSX) file. The output uses the same behavior as when you export to spreadsheet while viewing the web report. • Export to PDF: The report is saved as a PDF file. The output uses the same behavior as when you save to PDF while viewing the web report. |
| Save or Email Files | <p>Select one of the following to determine the delivery method for the output:</p> <ul style="list-style-type: none"> • Save Files (default): The output files are saved to the specified output folder. • Email Files: The output files are emailed to the specified recipients. The output files are not saved anywhere on the file system. • Save and Email Files: The output files are both saved and emailed. |

| Item | Description |
|-----------------|---|
| File Generation | <p>Select one of the following to determine whether the output is saved as a single file or multiple files:</p> <ul style="list-style-type: none"> • Multiple Output files (default): The results of each pass are saved as individual output files. For example, if the multipass settings result in 10 passes, then 10 output files are created (one file for each pass). • Single Output File: The results of each pass are collected into a single output file. For example, if the multipass settings result in 10 passes, then the results of all 10 passes are placed in a single output file. <p>If the output type is Excel, then each pass is a separate sheet in the Excel file. If the output type is PDF, then the PDF for each pass is combined into one large PDF file.</p> |
| File Name | <p>Specify how the output file (or files) should be named. You can do the following:</p> <ul style="list-style-type: none"> • You can use processing variables and/or Scheduler job variables to generate dynamic file names. • You can type a "hard-coded" file name. <p>If the task will generate multiple output files, then the file name (or the output folder path) must use a processing variable so that the output of each pass is unique. If the task will generate a single output file, then variables are not required.</p> <p>To use a processing variable, you can type the variable or you can click the pencil icon  to open a text editor. From the Insert Variable list, select the variable that you want to use.</p> <p>For example, you could set the file name to <code>Income Statement [Current_Value]</code>. If the report is being processed by region to multiple output files, this will generate file names such as <code>Income Statement West</code>, <code>Income Statement East</code>, and so on (where "East" and "West" are region names).</p> <p>NOTE: Processing variables and Scheduler variables use different syntax. Processing variables are enclosed in square brackets. Scheduler job variables are enclosed in curly brackets.</p> |

| Item | Description |
|------------|---|
| Sheet Name | <p>Specify how the sheet for each pass should be named. This property only applies when the processing type is Export to Excel. You can do the following:</p> <ul style="list-style-type: none"> You can use processing variables and/or Scheduler job variables to generate dynamic sheet names. You can type a "hard-coded" sheet name. <p>If the task will collect all of the output into a single spreadsheet file, then the sheet name must use a processing variable so that the output of each pass is unique. If the task will generate multiple output files, then variables are not required.</p> <p>To use a processing variable, you can type the variable or you can click the pencil icon  to open a text editor. From the Insert Variable list, select the variable that you want to use.</p> <p>For example, you could set the sheet name to <code>[Current_Value]</code>. If the report is being processed by region, this will generate sheet names such as <code>West</code>, <code>East</code>, and so on (where "East" and "West" are region names).</p> <p>NOTE: Processing variables and Scheduler variables use different syntax. Processing variables are enclosed in square brackets. Scheduler job variables are enclosed in curly brackets.</p> |

Export to Excel Settings

Complete the following properties if the processing type is Excel.

| Item | Description |
|------------------------|---|
| Include Column Headers | <p>Specifies whether column headers are included in the file output. By default this is set to On, which means column header text is included in the first row of the spreadsheet. Column grouping headers and multi-row headers are not included.</p> <p>If this option is set to Off, then column headers are omitted from the file output and the data starts in the first row of the spreadsheet.</p> |
| Include total row | <p>Specifies whether the total row is included in the file output. By default this is set to On, which means that the total row is included in the spreadsheet.</p> <p>If this option is set to Off, then the total row is omitted from the file output.</p> <p>NOTE: This option only applies when the web report being processed is a dynamic row report with the total row enabled. If the web report being processed uses a fixed row structure, then the total and subtotal rows defined in the fixed row structure are always included in the spreadsheet.</p> |

Export to PDF Settings

Complete the following properties if the processing type is PDF.


| Item | Description |
|-----------------|---|
| PDF Orientation | Select the orientation for the PDF, either Portrait or Landscape . Portrait is the default orientation. |
| Page Size | Select the page size for the PDF. You can choose from the following standard page sizes: A3 , A4 , A5 , Legal , Letter , or Tabloid . Letter is the default size. |

Output File Settings

Complete the following properties if file output is being saved. These settings do not apply if the output is email only.


| Item | Description |
|-----------|--|
| Output To | <p>Select one of the following:</p> <ul style="list-style-type: none">• Local File System (default): The output location is outside of Axiom, to a location on your local network share. The specific path is detailed in the Output Folder setting. Access to output files is not controlled by Axiom.• Axiom Repository: The output location is the Axiom file system, within the Reports Library. The specific path is detailed in the Output Folder setting. Access to output files is controlled by security access to the designated folder within Axiom. |



| Item | Description |
|---------------|--|
| Output Folder | <p>Specify the folder location for the file output. You can type a folder path, or you can click the folder icon to browse to the folder location. The browse dialog will display either your local file system or the Axiom file system, depending on what you selected for Output To.</p> <p>The output folder can be made dynamic as follows:</p> <ul style="list-style-type: none"> • If File Generation is set to Multiple Output Files, then processing variables can be used in the output folder path. For example, you can include <code>[Current_Value]</code> in the output folder path, and this will be replaced with the current multipass value. Processing variables are not valid in the output folder path if the task is configured to generate a single output file. • Scheduler job variables can be used in the output folder path. <p>NOTE: Processing variables and Scheduler variables use different syntax. Processing variables are enclosed in square brackets. Scheduler job variables are enclosed in curly brackets.</p> <p>Local file system</p> <p>The output folder location must be entered as a UNC path, and must be accessible by the Scheduler service user account (for on-premise systems) or the Axiom Cloud Integration Service (for cloud systems).</p> <p>The ability to save files to the specified location and access them after saving is controlled by local network security.</p> <p>Axiom repository</p> <p>The specified location in the Axiom file system must be within the Reports Library, and the location must use the full path (meaning: <code>\Axiom\Reports Library\...</code>). The ability to save files to the specified location and to create new folders (if necessary) depends on the Axiom security permissions for the user processing the file. Users can only create new folders if they have read/write permissions to the parent folder, and they can only create new files if they have read/write permissions to the target folder.</p> <p>Once the files are created within the Axiom file system, access to those files is dependent on the user's permissions to the output folder. Typically you should create the output folder in advance (or if you want to create output folders on-the-fly, create a parent folder to hold the output folders), and then set permissions for that folder as appropriate in Axiom security, so that the appropriate users will be able to access the files after they are created.</p> |

| Item | Description |
|------------------------|---|
| Remote Data Connection | <p>This option only applies when the file output is being saved to your local file system, and only for Axiom Cloud systems that are using remote data connections.</p> <p>Select the name of the remote data connection to use for the file processing operation. The designated remote data connection will be used to access the local file system and save output file(s) to the designated location.</p> <p>A remote data connection is required to save files locally from an Axiom Cloud system. For more information, see the section on remote data connections in the <i>Scheduler Guide</i>.</p> |
| Purge Setting | <p>This option only applies when the file output is being saved to the Axiom Repository.</p> <p>If you want the file output to be automatically deleted after a specified period of time, then click the pencil icon  to open the Choose Date dialog.</p> <ul style="list-style-type: none"> • No purge date (default): File output is not automatically deleted. • Static purge date: Select a specific date, after which the output will be deleted. • Relative purge date: Specify a number of days to keep the output after it has been generated. The output will be deleted after the specified number of days have passed. |

Email Settings

Complete the following properties if file output is being emailed. These settings do not apply if the output is saved only.



| Item | Description |
|------------------|---|
| Recipient column | <p>Optional. Specify a table column that holds the desired email recipients for each pass. This option only applies if File Generation is set to Multiple Output Files, so that each pass will be sent a separate email.</p> <p>You can type the name of a table column, or click the column button  to select a column from the multipass table or a lookup table. (You must select a multipass column first before you can use the column button to select a column.) For example, if the multipass column is Dept.VP, the recipient column might be Dept.VP.Email.</p> <p>The specified column can contain any of the following: email addresses, user login names, and/or role names. The column can contain multiple values separated by a semicolon. The recipients listed in the column will be used as the To address for the email (in addition to any recipients listed directly in the To field). If the column contains a user login name, that user's email address as defined in security will be used. If the column contains a role name, the email will be sent to all users in the role.</p> <p>To verify that the recipient column will resolve as you expect for each pass, you can click the Preview Multipass List button in the Multipass Data Settings section. The specified recipient column displays in this preview so that you can see the recipient column values associated with the multipass column values.</p> <p>NOTE: The recipient column must have a one-to-one relationship with the values in the specified multipass column.</p> |
| To | <p>Specify the To recipient(s) for the email. This is required if a recipient column is not specified. If a recipient column is specified, the recipients listed here will be added to the recipients listed in the column for each pass.</p> <p>You can type one or more email addresses, user login names, and/or role names. Separate multiple recipients with semicolons. If a user login name is listed, that user's email address as defined in security will be used. If a role name is listed, the email will be sent to all users in the role.</p> <p>NOTE: If File Generation is set to Multiple Output Files, the recipients in the To field will receive a separate email for each pass. The only way to dynamically send the emails to different recipients per pass is to use the Recipient Column option.</p> |
| CC | Optional. Specify the CC recipient(s) for the email. This field follows the same rules as the To field. |
| BCC | Optional. Specify the BCC recipient(s) for the email. This field follows the same rules as the To field. |



| Item | Description |
|--------------|--|
| From | <p>Select one of the following to specify the From address for the email:</p> <ul style="list-style-type: none"> • Current User: The email will be sent from the user who executes the Scheduler job. • System User: The email will be sent from the designated From user for Scheduler. This is the same value returned by the <code>{Scheduler.FromEmailAddress}</code> job variable. |
| Subject Line | <p>Enter the subject line for the email. Processing variables can be used in the subject line when File Generation is set to Multiple Output Files.</p> <p>To use a processing variable, you can type the variable or you can click the pencil icon  to open a text editor. From the Insert Variable list, select the variable that you want to use.</p> <p>For example, you could set the subject line to <code>Monthly report for [Current_Value]</code> in order to include the current pass value in the subject line.</p> |
| Body Text | <p>Enter the body text for the email. Processing variables can be used in the body text when File Generation is set to Multiple Output Files.</p> <p>To use a processing variable, you can type the variable or you can click the pencil icon  to open a text editor. From the Insert Variable list, select the variable that you want to use.</p> |

Scheduler job variables can be used in any of the email settings except the From setting.

► Multipass properties

The multipass properties at the bottom of the task determine how the report will be processed over a dimension.

| Item | Description |
|---------------------|--|
| Multipass Column | <p>Specify the column to use for multipass processing. You can type a Table.Column name, or click the column icon  to select the column from a dialog. You can select any column on a data or reference table, though typically processing is performed by a dimension such as Dept.Dept, or a grouping such as Dept.Region.</p> <p>The report will be processed once for each unique value in the specified column (except for any values excluded by the Source Filter). A filter is applied to the data query in the report so that the data is limited to the current pass value. For example, if you are processing by Dept.Dept, then the report will be processed once for each department, and the report data will be limited to only the data for that department.</p> <p>Keep in mind the difference between processing by a data table column such as GL2021.Dept, versus a dimension table column such as Dept.Dept. When processing by GL2021.Dept, the report will be processed by each department with data in the GL2021 table. When processing by Dept.Dept, the report will be processed by each department in the Dept table.</p> <p>To verify the list of values for processing, click the Preview Multipass List button to view the list of items. The first 100 values are shown, in the order they will be processed. If the task configuration includes a Recipient Column (in the email settings) or a Sort By column, these columns are also shown in the preview.</p> |
| Current Pass Header | <p>Optional. Define a header to display in the report output file. This option only applies if the processing type is Export to PDF.</p> <p>The current pass header should use processing variables to display information about the current pass. To use a processing variable, you can type the variable or you can click the pencil icon  to open a text editor. From the Insert Variable list, select the variable that you want to use.</p> <p>For example, you can define a header such as:</p> <pre>Processed by [MULTIPASS_COLUMN] [CURRENT_VALUE]</pre> <p>When processing by Dept.Dept, this would resolve such as Processed by Dept 22000</p> <p>By default, if the current pass header is left blank, then the PDF output will not include a header to indicate the current pass information. However, it is possible that the template used to create the report may have been designed with a dynamic header that will display this information.</p> |

| Item | Description |
|---------------|--|
| Sort By | <p>Optional. Specify one or more sort columns for the list of multipass values. You can type a Table.Column name, or click the column icon  to select the column from a dialog. You can also optionally specify <i>Asc</i> or <i>Desc</i> after the column name (ascending order is used if not specified). For example: Dept.Dept Desc. Separate multiple values with semicolons.</p> <p>By default, the values are sorted by the multipass column in ascending order. The Sort By field only needs to be used if you want the values to be sorted in descending order instead, or if you want the values sorted by a different column in the same table.</p> <p>The processing order is only relevant when File Generation is set to Single Output File, since it determines the order of each individual pass within the single file. When outputting to Multiple Output Files, the order is still used during processing but it has no useful impact on the outcome.</p> |
| Source Filter | <p>Optional. Specify a filter to limit the multipass list of items. You can type a filter, or you can click the filter icon  to use the Filter Wizard.</p> <p>When the multipass list of values is generated, any value that does not meet the source filter will be excluded from processing.</p> <p>By default, all values in the specified multipass column are processed if the source filter is left blank.</p> |

Scheduler job variables can be used in any of the multipass settings.

► Using processing variables

The following processing variables can be used in various settings within the Web Report Processing task, in order to dynamically change the setting using information for the current pass.

| Item | Description |
|----------------------|--|
| [CURRENT_VALUE] | <p>This variable returns the current multipass processing value. For example, if you are processing by Dept.Dept, and the current pass is for department 20000, the variable will be replaced by the value "20000" for this pass.</p> <p>This variable is typically used in settings such the file name, sheet name (when generating Excel output), and folder path.</p> |
| [CURRENT_PASSNUMBER] | <p>This variable returns the current pass number. For example, if the current pass is number 20 of 35 passes, the variable will be replaced by the value "20" for this pass.</p> |

| Item | Description |
|--------------------|---|
| [MULTIPASS_COLUMN] | <p>This variable returns the name of the multipass column. For example, if you are processing by Dept.Dept, the variable will be replaced by the value "Dept" for all passes.</p> <p>This variable could be used whenever you want to reference the name of the dimension processed. For example, instead of just referencing the current value in the file name, you might want to reference the column name and the value. A variable construction like [MULTIPASS_COLUMN] [CURRENT_VALUE] would resolve to "Dept 20000" when processing by Dept.Dept and the current pass is for department 20000.</p> |

Processing variables can only be used in certain settings, and sometimes only when the output is multiple files (versus a single file). See the documentation for each individual setting to see if processing variables are supported in that setting.

Scheduler tasks for database maintenance

Scheduler provides several built-in tasks that are intended for database maintenance. By default, these tasks are included in the System.IndexMaintenance job, which runs regularly to maintain your database. However, these tasks can also be added manually to jobs as needed to perform additional database maintenance.

The following database maintenance tasks are available:

- **Rebuild Database Indexes** task
- **Update Database Statistics** task
- **Update Indexes and Constraints** task


All of these tasks are predefined versions of the **Execute SQL Command** task. You can use the **Source Axiom Database** field to specify whether the task is executed against the system database or the audit database.

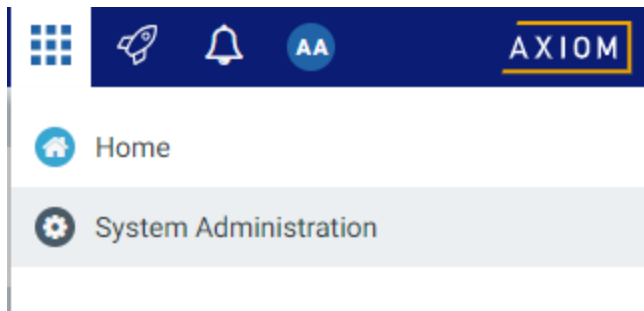
For the **SQL Command Text**, the actual SQL code used by each task is generated automatically by Axiom when the task is executed. This ensures that the tasks always use the most current SQL code for each task as defined by Axiom.

Web Scheduler

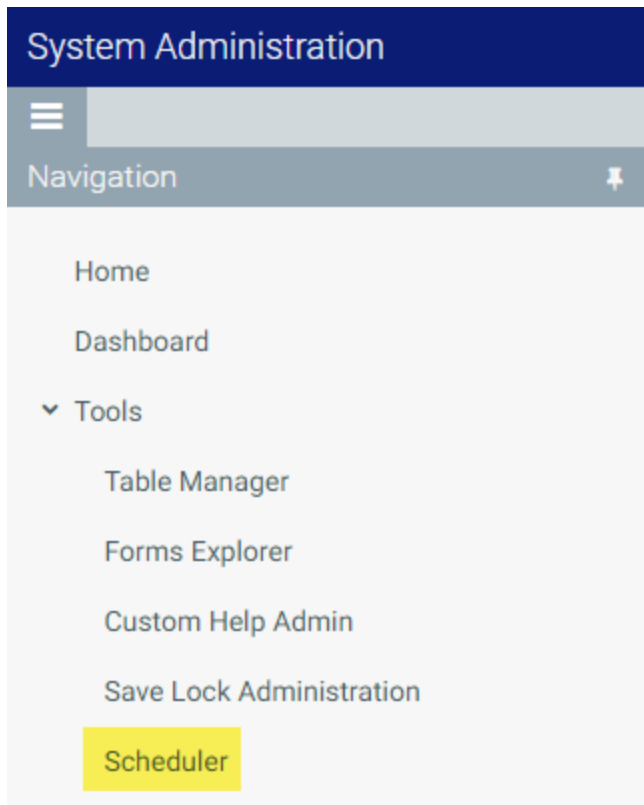
Although most Scheduler setup activities can only be performed in the Desktop Client, some job management can be performed in the Web Client. Using the "Web Scheduler", you can monitor and manage the job schedule, review job results, and process existing jobs on demand.

To access Scheduler in the Web Client:

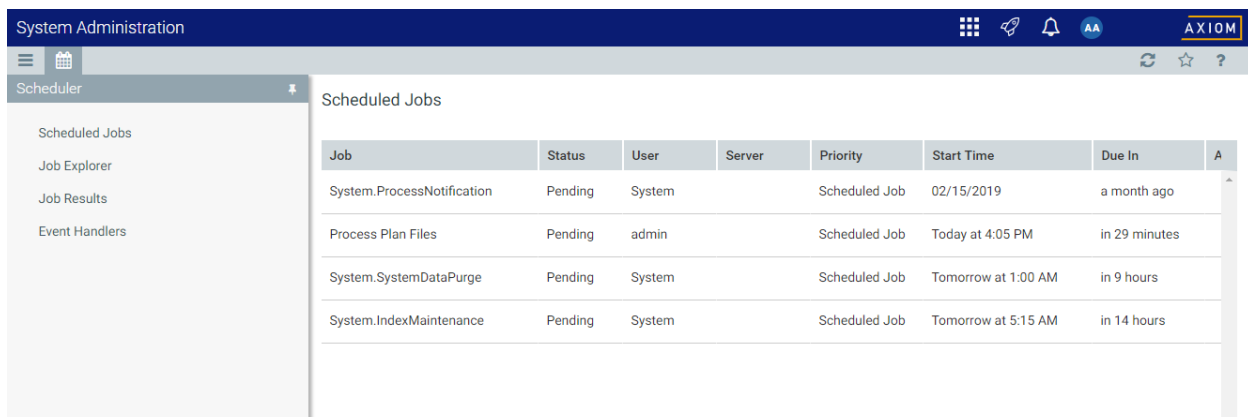
1. In the Web Client, click the menu icon  in the Global Navigation Bar. From the Area menu, select **System Administration**.



2. From the Navigation panel, select **Tools > Scheduler**.



When you access the Scheduler area, a Scheduler panel becomes available in the left side of the Task Bar. You can use this panel to change the current Scheduler view.



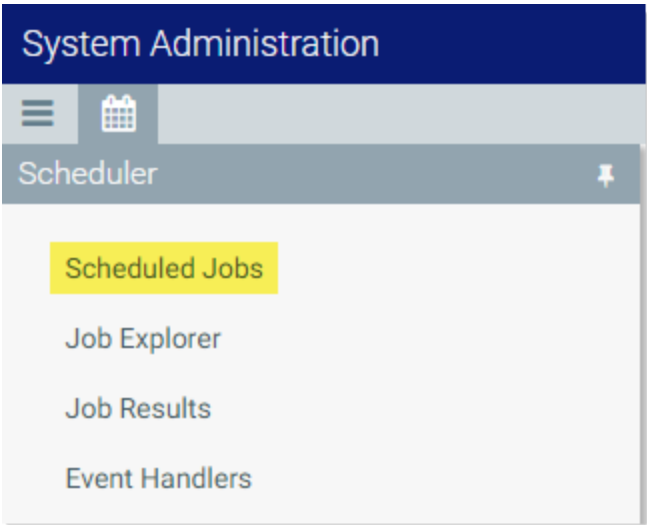
Example Scheduler area in Web Client

Managing the job schedule in the Web Client

In the Scheduler area of the Web Client, you can view the status of all jobs that are currently on the schedule. If necessary, you can view the job details, remove the job from the schedule, or run the job now.

To view the current job schedule:

- On the [Scheduler](#) page, from the **Scheduler** panel, select **Scheduled Jobs**.



The **Scheduled Jobs** grid displays all jobs that are scheduled to be processed, or are currently in process. This includes scheduled jobs, jobs executed manually via **Run Now**, and jobs that were triggered for execution via an event handler.

A screenshot of the 'System Administration' interface showing the 'Scheduled Jobs' grid. The grid has columns: Job, Status, User, Server, Priority, Start Time, Due In, and Actions. The 'Process Plan Files' job is highlighted in blue. The 'Actions' column for this job shows icons for view, edit, and run. The 'AXIOM' logo is visible in the top right corner.

| Job | Status | User | Server | Priority | Start Time | Due In | Actions |
|----------------------------|---------|--------|--------|---------------|---------------------|---------------|---------|
| File Processing | Working | admin | | Manual | Today at 3:37 PM | a minute ago | |
| System.ProcessNotification | Pending | System | | Scheduled Job | 02/15/2019 | a month ago | |
| Process Plan Files | Pending | admin | | Scheduled Job | Today at 4:05 PM | in 27 minutes | |
| System.SystemDataPurge | Pending | System | | Scheduled Job | Tomorrow at 1:00 AM | in 9 hours | |
| System.IndexMaintenance | Pending | System | | Scheduled Job | Tomorrow at 5:15 AM | in 14 hours | |

Example Scheduled Jobs grid

You can use the **Actions** column in the right side of the grid to perform any of the following actions on a job:



View the job properties and results.




Remove the job from the schedule.

IMPORTANT: If the job is on the schedule due to a scheduling rule, this action disables the scheduling rule and removes all future executions from the schedule as well. If you want future scheduled instances of the job to proceed, you must edit the job to re-enable the scheduling rule.



Run the job now.

This action places the job on the schedule for immediate execution (if another manual instance of the job is not already pending). The future scheduled instance of the job remains on the schedule.

To refresh the list, click the Refresh icon  in the Task Bar.

For each job on the schedule, you can view the following information:

| Item | Description |
|--------|---|
| Job | The name of the job. |
| User | <p>The user identity that the job will be run as. If the job is a system job, the user is System.</p> <p>This is typically the name of the user who placed the job on the schedule, but not always (for example, when using an event handler that is set to run as the job owner instead of as the requester).</p> |
| Status | Job status is either Pending (waiting to be executed) or Working (currently being executed). |
| Server | If a job is currently Working , then the server executing the job is listed here. Otherwise, this column is blank. |

| Item | Description |
|------------|---|
| Priority | <p>The priority category for the job:</p> <ol style="list-style-type: none"> 1. Manual: The job was executed manually. 2. Event Handler: The job was executed by a Scheduler event handler. 3. Scheduled Job: The scheduled instance of the job results from an active scheduling rule. 4. Subordinate Job: The job was generated as a subordinate job, from a currently executing job. <p>The priority category determines how jobs are evaluated for processing order, in conjunction with the job's Priority Elevation setting. Manual jobs are highest priority, and subordinate jobs are lowest priority. For more information, see Processing priority for scheduled jobs.</p> |
| Start Time | <p>The start time of the job. The job is eligible for immediate execution if the start time is now or passed. Jobs may not be executed right at the start time if no Scheduler threads are currently available to execute the job, or if other eligible jobs have higher priority.</p> <p>If the job is on the schedule due to a scheduling rule, the start time is based on the scheduling rule. If the job was manually executed via Run Now or triggered by an event handler, the start time is the time the execution was initiated.</p> |
| Due In | <p>The length of time until the job is due to be processed. For example, if the job is scheduled to run at noon and it is currently 11:50 AM, then the job is due to be run in 10 minutes.</p> <p>This column is intended to make it easy to see when a job will be run, without needing to calculate it based on the start time.</p> |


NOTE: If a job has a scheduling rule with a recurring schedule, only the first scheduled execution appears in the **Scheduled Jobs** list. For example, if you have a job that is scheduled to run once a month for a year, you will not see all twelve scheduled executions in the list—you will only see the first scheduled execution. Once that instance has been run, the scheduling rule is re-evaluated and the next scheduled execution appears in the list.

Viewing job results in the Web Client

In the Scheduler area of the Web Client, you can view the results of jobs that have been executed. For each job, you can see when it was run, and whether it completed successfully or had errors.


To view job results:


- On the [Scheduler page](#), from the **Scheduler** panel, select **Job Results**.




The **Job Results** grid shows a list of jobs that have been recently executed. The grid shows the following summary information:


- The name of the job, and the ID of the particular execution of that job
- The user identity the job was run as
- The status of the job, such as Success or Failure
- The Scheduler server that ran the job
- The start time and duration of the job

To refresh the list, click the Refresh icon  in the Task Bar.


| System Administration | | | | | | | |
|--|------------------------------------|-------|-------------|--------|------------------|---------------|---|
| Job Results | | | | | | | |
| <input checked="" type="checkbox"/> Hide system jobs | | | | | | | |
| ID | Job | User | Status | Server | Start Time | Duration | |
| 4585946 | File Processing | admin | In Progress | WHQA1 | Today at 3:25 PM | a few seconds |  |
| 4585937 | Process Plan Files | admin | Success | WHQA1 | Today at 3:17 PM | a few seconds | |
| 4585944 | Process Plan Files(4586885)[01][7] | admin | Success | WHQA1 | Today at 3:14 PM | a minute | |
| 4585943 | Process Plan Files(4586885)[01][6] | admin | Success | WHQA1 | Today at 3:14 PM | 3 minutes | |
| 4585942 | Process Plan Files(4586885)[01][5] | admin | Success | WHQA1 | Today at 3:14 PM | 3 minutes | |
| 4585941 | Process Plan Files(4586885)[01][4] | admin | Success | WHQA1 | Today at 3:11 PM | 3 minutes | |
| 4585940 | Process Plan Files(4586885)[01][3] | admin | Success | WHQA1 | Today at 3:11 PM | 3 minutes | |
| 4585939 | Process Plan Files(4586885)[01][2] | admin | Success | WHQA1 | Today at 3:11 PM | 3 minutes | |
| 4585938 | Process Plan Files(4586885)[01][1] | admin | Success | WHQA1 | Today at 3:11 PM | 3 minutes | |
| | | | | | | | 1 - 20 of 55 items |

Example Job Results grid

To view detailed results for a particular job execution, hover your cursor over the job and then click the View icon  in the far right column. This opens the job properties to the **Job Results** section, with the corresponding execution ID expanded. You can further expand the job results to see the specific tasks that were executed.

| System Administration | | | | | | | |
|---------------------------|-----------------------|--|---------|------------------|---------------|-------|---|
| File Processing Results | | | | | | | |
| ID | | | Result | Start Time | Duration | User | |
| ▼ 4585946 | | | Success | Today at 3:25 PM | a few seconds | admin | |
| ▼ Job: File Processing | Server: WHQA1 | | Success | Today at 3:25 PM | a few seconds | | |
| | Task: File Processing | | Success | Today at 3:25 PM | a few seconds | |  |
| ► 4585928 | | | Success | Today at 3:02 PM | a few seconds | admin | |

Example Job Results detail showing tasks executed

To view the detailed task results, hover your cursor over the task and then click the View icon  in the far right column. This opens a dialog to display the results for that task. For example, for a file processing task, the detailed results would contain information such as the processing type and the number of passes, and the output that was created at the end of the process.

Once you are viewing the Job Results section of the job properties, you can review all of the available job history as needed. Expand any execution ID to view the details for that particular execution.

TIP: You can also view job results by opening a job and viewing the job properties, which include the job results. In some cases it may be easier to open the job and review all of its results rather than trying to find the job within the overall job results. For more information, see [Viewing jobs and event handlers in the Web Client](#).

NOTE: Users with the **Scheduled Jobs User** security permission can only see job results for jobs that they executed. Administrators can see job results for all jobs.

► System job results

By default, system job results are hidden in the **Job Results** grid. System jobs such as the SMTP message delivery job may run frequently, and can easily fill up the result history, making it difficult to find results for user-initiated jobs.

If you want to view results for system jobs, you can do one of the following:

- Clear the **Hide system jobs** check box above the **Job Results** grid. The list immediately updates to include system jobs.
- Open the system job directly, and view its job results within the job. For example, you can go to the **Scheduled Jobs** page and double-click the **System.SystemDataPurge** job to view all results for that job.

► Job result availability

Job results are purged periodically to help optimize system performance. The availability of job results in your system depends on the configuration of the system job **System.PurgeSystemData**. This system job runs periodically to purge old data in your system, including old job results. By default, when this job is run, it purges job history older than 15 days.

The configuration of this system job can only be viewed and edited in the Desktop Client, and only by administrators. For more information, see [Configuring Scheduler system jobs](#).

Additionally, individual jobs can be configured to purge old results when the job is run. In the Web Client, you can view the job properties to see if this option is enabled, but you cannot edit the job properties. The option is displayed in the **General** section of the job, under **Job Results Cleanup**. For more information on viewing job properties, see [Viewing jobs and event handlers in the Web Client](#).


Running a job manually in the Web Client


In the Scheduler area of the Web Client, you can run a job manually as needed.

When using this approach, the job is run now. It is not possible to run a job manually and specify a future execution time. If you want to schedule a job for future execution, you must define a scheduling rule on the job, which can only be done in the Desktop Client. For more information, see [Defining scheduling rules for a job](#).

To run a Scheduler job manually:

1. On the [Scheduler page](#), from the **Scheduler** panel, select **Job Explorer**.



2. In the Job Explorer page, locate the job that you want to run. This page lists all jobs in the Scheduler Jobs Library that you have permission to access.
3. Hover your cursor over the job, then click the **Run Once** icon  in the far right column.

The job is added to the schedule with a start time of now, and is eligible for immediate execution (pending available Scheduler threads and any higher-priority jobs already in the queue). You are automatically taken to the **Scheduled Jobs** area of Scheduler, so that you can see the job on the schedule.

Running a job manually does not impact any scheduled executions of the job as determined by scheduling rules. For example, if a job is scheduled to be run at 10:00 PM tonight, and you run the job manually at 2:00 PM, the job will still be run as scheduled at 10:00 PM.

Viewing jobs and event handlers in the Web Client

In the Scheduler area of the Web Client, you can view Scheduler jobs and event handlers.


Viewing jobs


You can view any job in the Scheduler Jobs Library that you have permission to access.

Scheduler jobs are read-only in the Web Client. You can view the job properties to better understand the purpose of a particular job and the tasks that it performs. The Web Client does not support creating new jobs, editing existing jobs, or deleting jobs. If you need to perform any of those actions, you must use the Desktop Client. For more information, see [Scheduler Overview](#).


To view a Scheduler job:

1. On the [Scheduler page](#), from the **Scheduler** panel, select **Job Explorer**.



2. In the Job Explorer page, locate the job that you want to view. This page lists all jobs in the Scheduler Jobs Library that you have permission to access.
3. Hover your cursor over the job, then click the **View** icon  in the far right column.

The job opens, and the Scheduler panel updates to show the viewable sections of the job. You can switch between sections by selecting section names in the Scheduler panel. By default, the **General** section is shown.



Example job properties

All job properties are defined in the Desktop Client. The following is a brief overview of the job properties shown in the Web Client.

| Section | Description | More Information |
|-----------|---|-------------------------------------|
| General | General job properties that impact the job's processing priority and processing behavior. | Job properties |
| Variables | Variables used by the job. <ul style="list-style-type: none"> If the job has defined variables, those variables display in the Job Variables section at the top of the page. Most likely, the tasks in the job are configured to use these variables. This typically means that the job is designed to be run using an event handler, and the necessary variable values will be passed to the job when it is triggered. The System Variables section displays the job's values for various system-defined variables. This section can help you understand who the owner of the job is, and how other system variables will resolve for the job. | Using job variables |

| Section | Description | More Information |
|------------------|--|---|
| Scheduling Rules | <p>Scheduling rules to schedule jobs for future execution. If the job has an active scheduling rule, the job will be executed according to the rule (one time or recurring, depending on how the rule is configured).</p> <ul style="list-style-type: none"> • Day of Week, Hours, and Minutes specify when the job will be executed within the start / end range of the rule. An asterisk in any of these fields means "all"—for example, if Hours is set to * then the job is run every hour. • Starting On and Ending On determine the start / end range of the rule. If they are blank, then the rule has no start or end date. | Defining scheduling rules for a job |
| Event Handlers | If the job is designed to be run using an event handler, the event handler name is listed here. The Execute As property determines whether the job is run as the requester or the job owner when it is triggered for execution. | Viewing event handlers |
| Notifications | Notification settings for the job. The job can be configured to send email notifications when the job completes, or only when the job has errors. Variables can be used to determine the notification recipients. | Setting up notifications for jobs |
| Tasks | Tasks to be executed by the job, listed by name and task type. No other task properties are available in the Web Client. If you want to see more information about the task, you must view the job in the Desktop Client. | Scheduler Task Reference |
| Results | Detailed results of the previous job executions. Results are organized by execution ID and displayed in execution order (the most recent listed first). | Viewing job results in the Web Client |

► Viewing event handlers


You can view the event handlers that are defined in the system. Event handlers are used to trigger Scheduler jobs based on an event.

For example, an Axiom form can have a Button component that is configured with the RunEvent command. When a user clicks the button, the specified event name is passed to Scheduler, and any jobs associated with that event are triggered to run. Variable values can also be passed from the form to the Scheduler job as part of this process.

Event handlers are read-only in the Web Client. If you need to create, edit, or delete an event handler, this can only be done in the Desktop Client. For more information, see [Managing event handlers](#).

To view Scheduler event handlers:

- On the [Scheduler page](#), from the **Scheduler** panel, select **Event Handlers**.



The **Event Handlers** grid lists all of the event handlers as follows:

- **Event Name:** Name of the event handler. This is the name used in features such as RunEvent to trigger execution of a Scheduler job.
- **Job:** Name of the job where the event name is used. When the event is raised by a feature such as RunEvent, this job will be executed.
- **User:** The user identity that will be used to execute jobs triggered by the event handler. If the event handler is configured to run as the requester, then **Requester** is listed here. If the event handler is configured to run as the owner, then the owner name is listed here (either a specific user name, or **System**).

Index

A

- Active Directory Import task 73
- Active Directory synchronization
 - import task 73
- Axiom Cloud 67

C

- Collect Worksheets task 77
- Copy On Demand Plan Files task 79
- Create Plan Files task 86

E

- Echo task 88
- email delivery 12
- emailing
 - setup 10
- event handlers 142
 - system event handlers 66
 - user-defined 31
- Execute Command Adapter task 89
- Execute SQL Command task 90
- Export ETL Package task 91

F

- File Processing task 91

I

- Import ETL Package task 96

J

- job result history (Scheduler)
 - purging 63
 - viewing 62
- jobs
 - about 3

- creating 20
- deleting 23
- editing 22
- executing using RunEvent 46
- history 62
- removing from schedule 60, 135
- rescheduling an execution 60
- results 138
 - purging 63
 - viewing 62
- running manually 46, 135, 141
- schedule 60, 135
- triggering other jobs 44
- variables 33
- viewing 142

N

- notifications
 - Scheduler 28

P

- parallel processing 36
- Process Document List task 107
- process management
 - scheduling processes 120
- Process Plan Files task 97
- Process Template List task 112
- Purge System Data task 115
- purging job results 63

R

- Raise Event task 116
- Rebuild Database Indexes task 133
- remote data connections
 - creating 67

Run Scheduler job task 117

RunEvent

 setting up in Scheduler 46

S

Scheduler 3

 about 3

 conditional processing 42

 dialog overview 6

 email delivery, configuring 12

 email notifications 10, 28

 enabling services 9

 event handlers 66

 hosted connections 67

 iterative processing 38

 job result history

 purging 10, 63

 viewing 62

 job variables 33

 jobs

 about 3

 managing 20

 running manually 46

 schedule 60

 scheduling rules 23

 settings 52

 parallel processing 36

 processing priority 5

 servers, managing 63

 setup 10

 system jobs 5

 tasks 70

 about 3, 71

 conditional processing 42

 managing 20

 properties 71

 renaming 58

scheduling rules

 adding 24

 defining 26

 deleting 27

SMTP 12

SMTP Message Delivery task 119

Start Process task 120

system jobs 5

System.SMTPMessageDelivery 12

T

tasks 70

 about 3

U

Update Database Statistics task 133

Update Indexes and Constraints task 133

V

variables

 job variables 33

W

Web Report Processing task 121

web reports

 production reporting 121

Web Scheduler 134

 event handlers 142

 jobs

 results 138

 running manually 135, 141

 schedule 135

 viewing 142