



THE
POWER
TO KNOW.

SAS[®] Factory Miner 14.1

Administration and Configuration

The correct bibliographic citation for this manual is as follows: SAS Institute Inc. 2015. *SAS® Factory Miner 14.1: Administration and Configuration*. Cary, NC: SAS Institute Inc.

SAS® Factory Miner 14.1: Administration and Configuration

Copyright © 2015, SAS Institute Inc., Cary, NC, USA

All rights reserved. Produced in the United States of America.

For a hard-copy book: No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, or otherwise, without the prior written permission of the publisher, SAS Institute Inc.

For a web download or e-book: Your use of this publication shall be governed by the terms established by the vendor at the time you acquire this publication.

The scanning, uploading, and distribution of this book via the Internet or any other means without the permission of the publisher is illegal and punishable by law. Please purchase only authorized electronic editions and do not participate in or encourage electronic piracy of copyrighted materials. Your support of others' rights is appreciated.

U.S. Government License Rights; Restricted Rights: The Software and its documentation is commercial computer software developed at private expense and is provided with RESTRICTED RIGHTS to the United States Government. Use, duplication or disclosure of the Software by the United States Government is subject to the license terms of this Agreement pursuant to, as applicable, FAR 12.212, DFAR 227.7202-1(a), DFAR 227.7202-3(a) and DFAR 227.7202-4 and, to the extent required under U.S. federal law, the minimum restricted rights as set out in FAR 52.227-19 (DEC 2007). If FAR 52.227-19 is applicable, this provision serves as notice under clause (c) thereof and no other notice is required to be affixed to the Software or documentation. The Government's rights in Software and documentation shall be only those set forth in this Agreement.

SAS Institute Inc., SAS Campus Drive, Cary, North Carolina 27513-2414.

September 2015

SAS® and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration.

Other brand and product names are trademarks of their respective companies.

Contents

<i>SAS Factory Miner 14.1: Administration and Configuration</i>	v
Chapter 1 • Introduction to SAS Factory Miner	1
Overview of SAS Factory Miner	1
Intended Audience	1
SAS Information Online	2
Chapter 2 • SAS Factory Miner System Requirements	3
System Requirements	3
Chapter 3 • Installing SAS Factory Miner	5
Installation Overview	5
Pre-Installation Tasks	6
Install and Configure SAS Factory Miner	9
Migrate SAS Factory Miner	10
Verify the Installation	11
Chapter 4 • Configuration and Administration Tasks	13
Specifying Security Permissions for Users and Groups	13
Performance and Data Governance Settings	15
Define the SAS Factory Miner Data	18
Project Directory and File Information	19

SAS Factory Miner 14.1: Administration and Configuration

Audience

This book is intended for system administrators who are installing SAS Factory Miner.

Chapter 1

Introduction to SAS Factory Miner

Overview of SAS Factory Miner	1
Intended Audience	1
SAS Information Online	2
Current Documentation	2
SAS Notes	2
SAS Technical Support Services	2
Product Help	2

Overview of SAS Factory Miner

SAS Factory Miner is a web-based application that enables you to rapidly build and compare several models within and across data segments. Before creating any models, you define a collection of model templates. A template specifies both the model used and the properties used to configure the model. Then, you define a segmentation variable that partitions the data. After you specify which model templates to apply, SAS Factory Miner applies those templates within each segment and chooses a champion model. Models can then be exported to SAS Model Manager.

Intended Audience

This administrator's guide describes the recommended system architecture and requirements, installation and post-installation tasks, ongoing system administration, and troubleshooting. This book is for administrators who need to install, configure, and optimize SAS Factory Miner. No programming experience, SAS or otherwise, is necessary to administer SAS Factory Miner.

The administrative duties are divided into two roles: a system administrator and a SAS Factory Miner administrator. These do not need to be the same person.

- The system administrator must be familiar with the hardware, operating system, file system, installed applications, maintenance, and operation of the computer system that will house SAS Factory Miner.
- The SAS Factory Miner administrator must be familiar with the installation, configuration, and maintenance of all parts of the SAS system, including the SAS

Metadata Server. One such duty is the maintenance of the metadata server, which regulates traffic among servers, users, and other global resources required by SAS Factory Miner.

SAS Information Online

Current Documentation

For the most current installation and configuration information for SAS Factory Miner, see <http://support.sas.com/documentation/onlinedoc>.

SAS Notes

For additional information and support fixes, you should check the SAS Notes that are available on the SAS Technical Support website. You can search for available SAS Notes for SAS Factory Miner at <http://support.sas.com/software/index.html>.

SAS Technical Support Services

The SAS Technical Support staff is available to respond to problems and answer technical questions.

Product Help

For information about how to operate your software and its features, select the **Help** menu item from within SAS Factory Miner.

Chapter 2

SAS Factory Miner System Requirements

System Requirements	3
Supported Operating Systems	3
Segment Variable Name Restrictions	3
Supported Internet Browsers	3
Hard Disk Space	3
RAM	4
CPU	4

System Requirements

Supported Operating Systems

The SAS Factory Miner server, middle tier, and workstation components are supported on the following operating systems:

- 64-bit Windows systems, including Windows 7 and Windows Server 2008
- Linux for x64

Segment Variable Name Restrictions

Due to how project folders are created, segment variable values cannot contain any special characters. That is, only alphanumeric characters, the underscore, and the hyphen are allowed. Users cannot register models when special characters are encountered.

Supported Internet Browsers

The list of supported browsers can be found at <http://support.sas.com/resources/thirdpartysupport/v94/browsers.html>.

Hard Disk Space

The SAS Factory Miner server installation uses approximately 30 gigabytes of hard disk space. In addition to the amount of disk space required to install SAS Factory Miner, you need adequate disk space to store the input data and the working data generated by SAS Factory Miner.

RAM

The SAS Factory Miner server requires 6 gigabytes of memory, plus 2 gigabytes per concurrent user. The mid-tier requires 8 gigabytes of memory. Therefore, a single machine installation would require a minimum of 16 gigabytes of memory.

CPU

The number of CPUs or processor cores needed to optimize performance for SAS Factory Miner on your system is difficult to specify. Factors such as the size of your data, I/O subsystems, the number of concurrent users, the CPU make and model, and your network infrastructure can affect the performance of any system.

Chapter 3

Installing SAS Factory Miner

Installation Overview	5
Pre-Installation Tasks	6
Overview	6
Deployment Plan	6
Create a SAS Software Depot	7
Additional Documentation	7
System Preparations	7
Designate Ports	8
The SAS Environment File	8
Third-Party Software for Windows	9
Determine the Project Directory	9
Install and Configure SAS Factory Miner	9
Migrate SAS Factory Miner	10
Verify the Installation	11

Installation Overview

To install SAS Factory Miner, use the SAS Deployment Wizard and follow the basic process described in the *SAS Enterprise Miner: Administration and Configuration Guide*. Some steps are specific to SAS Factory Miner and are covered in this chapter.

To install SAS Factory Miner, you must use a deployment plan file. A deployment plan is an XML file that is used as input to the SAS Deployment Wizard. SAS Factory Miner uses a deployment plan file. You can use either the plan that is included in the SAS Software Depot or a custom plan file. See [Deployment Plan on page 6](#) for more information. *The SAS 9.4 Intelligence Platform: Installation and Configuration Guide* provides detailed information about deployment plans and pre-installation tasks.

Pre-Installation Tasks

Overview

Before you begin your installation of SAS Factory Miner, you must complete the following tasks:

1. Create the Deployment Plan file.
2. Create a SAS Software Depot.
3. Review additional documentation.
4. Prepare your computer systems.
5. Designate ports.
6. Define the SAS Environment File Path.
7. Install required third-party software.
8. Determine where project files are stored.

Deployment Plan

When you and your SAS representative initially plan the deployment of SAS Factory Miner, your SAS representative will record your decisions about which software you will deploy and on which hardware that software will be installed. This information is used to generate a deployment plan file.

A deployment plan is an XML file that is used as input to the SAS Deployment Wizard. There are two types of deployment plans: standard and custom. A standard deployment plan describes a common configuration. Standard deployment plans are stored by default in the SAS Software Depot. A custom deployment plan is created by a SAS representative specifically for a site.

If your environment requires a custom plan, SAS sends you an email message with either a set of files or a ZIP file that has a set of files. Those files include the following:

- The Deployment file (plan.xml) — This is one of the key files in the SAS Project Directory that is used throughout the rest of the deployment process to customize your installation and configuration experience. The plan.xml file serves as input to the SAS Deployment Wizard. Copy this set of files to the plan_files folder on the top level of the SAS Software Depot.
- The Pre-Installation Checklist (checklist.pdf or checklist.rtf) — This is where you record information such as the credentials for special SAS users, machine names, ports, and the installation locations of third-party products relevant to your deployment.
- The Index file (index.html) — This file provides some high-level guidance through the remainder of the process. The index.html file includes a reference to a pre-installation checklist, which is for the information that you need to know for the remainder of the process. Information that you enter in the checklist is requested in later stages of the installation process. It is important that you record the information for later use to ensure consistency and to document decisions that you made.

The SAS 9.4 Intelligence Platform: Installation and Configuration Guide provides detailed information about deployment plans and pre-installation tasks.

Create a SAS Software Depot

When you download a SAS order, the SAS Download Manager automatically creates a SAS Software Depot on your system. For more details about creating your software depot, see “Creating a SAS Software Depot” in the *SAS 9.4 Intelligence Platform: Installation and Configuration Guide*.

Additional Documentation

The *SAS 9.4 Intelligence Platform: Installation and Configuration Guide* provides detailed information related to all installations that include all or part of the SAS 9.4 Intelligence Platform. SAS Factory Miner is bundled with several other SAS products that are not covered in detail by this guide. It is important that you become familiar with the *SAS 9.4 Intelligence Platform: Installation and Configuration Guide*.

You should also familiarize yourself with the following:

- Quick Start Guides and System Requirements — <http://support.sas.com/documentation/installcenter/94/index.html>. Select your operating system from the list at the bottom of the page.
- Alert-level SAS Notes — <http://support.sas.com/kb/?ct=51000&col=supprdr>. Click **Search Problem Notes** in the left navigation pane.

System Preparations

As a pre-installation task, you must designate or create a domain account that acts as the SAS Installer. This account requires administrative operating system privileges on each machine that hosts the SAS Enterprise Miner server, a SAS Workspace server, or SAS Factory Miner project files.

On a Windows system, this account also requires the permissions **log on as a batch job** and **Access this computer from the network**. The logon permission permits the user to be logged on by means of a batch queue. You can configure this security setting by opening the appropriate policy and adding each account.

Generally, it is recommended that you create a user group, and give the group the **Log on as a batch job**, and, if necessary, the **Access this computer from the network** permission. Then, add each account to that group. For UNIX, the SAS Installer account should not be the root account.

The **Access this computer from the network** permission is also required in order to connect to SAS servers. Typically, this right is already granted to the Windows group Everyone. To confirm, check the Windows local policy settings.

The SAS Factory Miner server uses both regular user accounts and product administrative user accounts. You can create regular user accounts for SAS Factory Miner as a post-installation task. For more information, see [Assign Users and Groups on page 14](#).

Note: A product administrator account is not the same as a SAS internal administrator account, such as the SAS Administrator (sasadm@saspw).

Designate Ports

The communications between the components of SAS Factory Miner require a number of ports that must be designated during the installation process. Consult the pre-installation checklist that is provided by your SAS representative for a complete list of ports that you must designate. The *SAS 9.4 Intelligence Platform: Installation and Configuration Guide* has more information about port requirements and selection in “Setting Up Users, Groups, and Ports.”

SAS Factory Miner requires its own workspace server and pooled workspace server that are configured in separate runs of the SAS Deployment Wizard; therefore, you need to designate unique ports for each server. If you specify the same ports in your second run of the SAS Deployment Wizard as you did in the first run, then you are unable to start the servers.

Note: Refer to the Pre-Installation Checklist for guidance in selecting appropriate port numbers. Determine and record the appropriate port numbers for the additional SAS Factory Miner workspace server and SAS Factory Miner pooled workspace server.

The SAS Environment File

A SAS environment file defines the available set of SAS environments for the SAS Factory Miner server’s client applications, and is generated during the configuration of the SAS Web Infrastructure Platform. This step is carried out by the SAS Deployment manager when you install the SAS Factory Miner server. The SAS Logon Manager includes a servlet that provides default information for the initial deployment.

When you have validated that your client applications work successfully with a deployment, it is recommended that you deploy the `sas-environment.xml` file to an HTTP server. This step ensures that you can customize the `sas-environment.xml` file to specify the name that you want to use and to account for the IT topology at your site. Your site might require that application clients interact with separate development, test, and production environments. Or you might choose to have separate SAS deployments to support distinct business units. In either scenario, when multiple environments are required, you can customize and deploy the `sas-environment.xml` file as needed.

Before starting the SAS Factory Miner server installation, you must determine the URL that is needed to access the default installation `sas-environment.xml`. The URL to access the default environment definition would be `http://<host>:<port>/sas/sas-environment.xml`. For example, if you are running a SAS Application Server on your desktop, you can deploy the default `sas-environment.xml` file at `http://localhost:80/sas/sas-environment.xml`.

For the SAS Factory Miner server, the default URL would resemble the following: `http://myserver.mynet.mycompany.com:80/sas/sas-environment.xml` on Windows and `http://myserver.mynet.mycompany.com:7980/sas/sas-environment.xml` on Linux.

For more details about the SAS environment file and how to configure SAS environments pre- or post-installation see the *SAS 9.4 Intelligence Platform: Middle-Tier Administration Guide*.

Note: An HTTP server is included with most installation plans. If you configured the SAS Web Server, then you do not need to deploy the SAS environment file.

Third-Party Software for Windows

The SAS Factory Miner server installation requires Windows PowerShell 2.0. You can determine the version of PowerShell by opening the Windows PowerShell command prompt and entering `$PSVersionTable`.

When installing SAS Factory Miner, you must provide valid paths to the required third-party software, or the SAS Deployment Wizard will not let you continue with the SAS Factory Miner server installation. For information about required versions of third-party software, see the Third-Party Software for SAS 9.4 available from the SAS Support Install Center, located at <http://support.sas.com/resources/thirdpartysupport/>.

Determine the Project Directory

Before installing SAS Factory Miner, you should determine where project files for SAS Factory Miner will reside. In the SAS Deployment Wizard, specify your chosen location on the SAS Factory Miner Server Configuration Content page in the **Project Directory** field. The default project directory is `SASConfig/Lev [n] /AppData/FactoryMinerServer`.

Your project directory should be on a network location if the pooled workspace server is load balanced.

If you deploy SAS Factory Miner with a grid, then the project directory must be accessible by all machines on the grid.

For Linux deployments, the **Factory Miner: User** role requires Read and Write privileges. Because SAS Factory Miner uses the Pooled Workspace Server to access the project directory, the Pooled Workspace Server requires Write privileges.

Project files are not removed when you uninstall SAS Factory Miner. If you need to remove this data, you must manually delete the project directory after uninstalling SAS Factory Miner.

See [Project Directory on page 19](#) for more information.

Install and Configure SAS Factory Miner

For the best overall explanation of the installation and configuration of the numerous combinations of products and environments, see the *SAS 9.4 Intelligence Platform: Installation and Configuration Guide*.

Use the SAS Deployment Wizard to install and configure your software. For more information about the wizard, see the *SAS Deployment Wizard User's Guide*.

Here is the information that is specific to installing SAS Factory Miner:

1. Start the SAS Deployment Wizard from your SAS Software Depot. For example, on a Windows system, double-click on the setup.exe file located in your SAS Software Depot folder.
2. In the Specify Deployment Plan step, select your deployment plan. You can use either the default plan file provided by the SAS Deployment Wizard or your own custom plan file.
3. In the Select Deployment Step and Products to Install step, select **Step 1**. The exact name of this step varies based on the number of machines in your deployment.

To install SAS Factory Miner, you must run the SAS Deployment Wizard twice. For a single-machine installation, you run the SAS Deployment Wizard twice on the same machine. This machine is the server, middle tier, and client. SAS Factory Miner requires a workspace server and a pooled workspace server configured with specific settings that are performed in the second run.

When your server tier and middle tier are different machines, you must run the SAS Deployment Wizard twice on the SAS Application Server in order to install and configure SAS Factory Miner. SAS Factory Miner requires a workspace server and a pooled workspace server configured with specific settings that are performed in the second run.

4. In the Select Configuration Prompting Level step, **Typical** is selected by default. You can also specify **Custom**, but not **Express**. This is required to enable token-based authentication.
5. Proceed with all other installation steps as described in the *SAS 9.4 Intelligence Platform: Installation and Configuration Guide*. You must complete the installation before moving on to the next step.
6. After the initial installation is finished, start the SAS Deployment Wizard for a second time and proceed to the Select Deployment Step and Products to Install step. This time, select **Step 2**. The exact name of this step varies based on the number of machines in your deployment.
7. In the Select Deployment Type step, select **Perform a Planned Deployment** and enable only **Configure SAS Software**.
8. During the second pass of the SAS Deployment Wizard, you are prompted for an external account for the SAS Spawned Servers Account. You can specify the same account information used in the first pass of the SAS Deployment Wizard.
9. In the Token-based authentication step, enable **Use SAS token authentication**.
10. In the SAS Application Server: Server Context step, enter a name in the **SAS Application Server Context Name** field that you can identify with SAS Factory Miner. For example, you might use **SASAppFM**.
11. In the SAS Factory Miner Context Configuration step, the default value should be the **Application Server Context Name** that you entered in the prior step. Do not change this value.
12. In the SAS Factory Miner Server Configuration Content step, specify the project directory that you identified in the pre-installation tasks.
13. Proceed with all other configuration steps until the SAS Deployment Wizard is finished.
14. Restart the SAS Object Spawner to ensure that the new servers are identified.

Migrate SAS Factory Miner

To migrate a SAS Factory Miner installation, you only need to run the SAS Deployment Wizard once. Although the deployment steps in your plan file indicate a second pass, it is not necessary.








When you migrate SAS Factory Miner, you must use the **Custom** path and manually select **SAS Factory Miner Configuration** on the Select Products to Configure page. For

single machine deployments, you must manually select the **SAS Factory Miner Mid-Tier** and **SAS Factory Miner Services API** middle tier components.

Project data is not moved during migration. You must manually move project data to your new location.

Verify the Installation

To verify that SAS Factory Miner was successfully installed, complete the following steps:

1. Complete the steps in [Define Additional Users and Groups on page 13](#) to provide one or more users with access to SAS Factory Miner.
2. Complete the steps to [Define a Data Library on page 18](#) in order to have accessible data.
3. Open a supported web browser and enter the URL to your SAS Factory Miner server. The standard URL is of the form *http://myserver.mycompany.com/SASFactoryMiner* for Windows and *http://myserver.mycompany.com:7980/SASFactoryMiner* for Linux.
4. Sign in to SAS Factory Miner with one of the users that you assigned in Step 1.
5. Click the menu button in the upper left corner and select **Data Sources** in the **SAS Factory Miner** section.
6. Click the new data source button  in the upper right corner. Navigate to the Purchases data source that you defined in the Step 2. Select that data source and click **Open**.
7. Click the menu button  in the upper left corner and select **Projects**.
8. Click the new project button  in the upper right corner. Enter a **Name** and **Data source** for this project. Click **Save**.
9. On the **Data** page, select the variable **Purchase**. Click . Set the **Role** of this variable to **Target**. Click **Save**.
Detailed instructions for this task are provided in the Help, which is accessible with the Help button .
10. On the **Data** page, select the variable **Region**. Click . Set the **Role** of this variable to **Segment**. Click **Save**.
Detailed instructions for this task are provided in the Help, which is accessible with the Help button .
11. Click **Build Profile**. Review the data segments created by SAS Factory Miner.
12. Click **Model Templates** to view the model templates included in this project.
13. Click **Run**. SAS Factory Miner takes you to the **Results** page. Review the results to ensure that there are no errors caused by an incomplete or faulty installation. If your trial data set does not model well, you might encounter errors in the model building step unrelated to installation errors.

A complete example using the Purchases data set can be found in the *SAS Factory Miner: User's Guide*.

Chapter 4

Configuration and Administration Tasks

Specifying Security Permissions for Users and Groups	13
Initial User and Role	13
Dependence on SAS Enterprise Miner	13
Assign Users and Groups	14
Integration with SAS Model Manager	14
Performance and Data Governance Settings	15
Model Creation Properties	15
Performance Tuning Recommendations	15
Stopped Models	17
Maximum Number of Threads	17
Exhausting the Pooled Workspace Server	18
Define the SAS Factory Miner Data	18
Define a New Library	18
Register a Data Set	19
Project Directory and File Information	19
Project Directory	19
Migration Information	20

Specifying Security Permissions for Users and Groups

Initial User and Role

When you install and configure SAS Factory Miner, standard user and role definitions are added. You must create external users and add them to the **Factory Miner: User** role to provide users with access to SAS Factory Miner.

Dependence on SAS Enterprise Miner

SAS Factory Miner has a dependency on SAS Enterprise Miner. Therefore, all SAS Enterprise Miner required user privileges are also required by SAS Factory Miner. This includes granting users the **Log on as a Batch Job** permission.

Assign Users and Groups

Complete the following steps to add users and groups to the **Factory Miner: User** role:

1. In SAS Management Console, select the **User Manager** plug-in.
2. Right-click the **Factory Miner: User** role and select **Properties**.
3. On the **Members** tab, move the desired users and groups from the **Available Identities** field to the **Current Members** field.
4. Click **OK**.

Integration with SAS Model Manager

SAS Factory Miner enables you to automatically register models with SAS Model Manager. In order to successfully register a model with SAS Model Manager, users must have the proper permissions to access SAS Model Manager.

Complete the following steps to ensure that users can register models with SAS Model Manager:

1. In SAS Management Console, select the **User Manager** plug-in.
2. Right-click the **Model Manager Advanced Users** group and select **Properties**.

Note: Users in the **Model Manager Administrator Users** group can register models with SAS Model Manager. Assign users to the appropriate group based on their needs in SAS Model Manager. Users in the **Model Manager Administrator Users** group do not also need to be in the **Model Manager Advanced Users** group.

3. On the **Members** tab, move the desired users and groups from the **Available Identities** field to the **Current Members** field.
4. Click **OK**.

Decision Builder supports only models with a score code type of DATA step or DS2.

For projects that are shared between SAS Factory Miner and SAS Model Manager, you must log on via the SAS Visual Analytics Hub using an external account. Failure to do so causes your projects to become out of sync between products.

Your SAS Factory Miner project folder structure must follow the same conventions as your SAS Model Manager folder structure. The subfolders in the Shared Data directory must have the same path structure and identical user permissions must be assigned.

If you are using SAS Factory Miner 14.1, then you must also use SAS Model Manager 14.1. For more information about SAS Model Manager, see the *SAS Model Manager 14.1: Administrator's Guide*.

Performance and Data Governance Settings

Model Creation Properties

In SAS Management Console, the **Factory Miner Services API** application management plug-in contains several settings that impact how models are built and champion models are selected.

To access these settings, complete the following steps:

1. Open SAS Management Console.
2. Expand the **Application Management** ⇨ **Configuration Manager** ⇨ **SAS Application Infrastructure** folder, right-click **Factory Miner Services API**, and select **Properties**.
3. Click the **Settings** tab.
4. You can edit the settings in the **Policies**, **Misc**, **Sample Data**, **Partition Data**, and **Rules** groups. Some of the settings defined here include sample size, whether data is partitioned, champion model selection criterion, and minimum event rate. Your business needs and analytical requirements determine the most appropriate values for these settings.

Performance Tuning Recommendations

Large Data Sets

Data sets that are larger than 10 gigabytes might require a key column in order to build a profile. If building the profile fails, users are notified only that there was an error building the profile. Users might not know that their data set requires a key variable.

Job Execution Services

The default Job Execution Services time out value might be too small to prevent time-outs when running a model training job. To increase this value, complete the following steps:

1. Open SAS Management Console.
2. Expand the **Environment Manager** ⇨ **Server Manager** ⇨ **SASAppFM** plug-in.
3. Right-click **SASAppFM — Logical Pooled Workspace Server** and select **Properties**.
4. Select the **Load Balancing** tab.
5. Change the value of **Availability timeout (sec)** to **1800**. This increases the time-out value to 30 minutes so that Job Execution Services does not time out while waiting for a pooled workspace server to run a model training job.
6. Click **OK**.

Model Training Jobs

For single-tier installations, it is recommended that you run only one model training job at a time. For grid installations, it is recommended that you run only one model training job at a time on each node. The steps to ensure this are the same.

1. Open SAS Management Console.
2. Expand the **Environment Manager** ⇒ **Server Manager** ⇒ **SASAppFM** ⇒ **SASAppFM — Logical Workspace Server** plug-in.
3. Right-click **SASAppFM — Pooled Workspace Server** and select **Properties**.
4. Select the **Options** tab and click **Advanced Options**.
5. Change the value of **Server process maximum** to **1**. Click **OK**.
6. Click **OK**.

Memory Settings

It is recommended that the MEMSIZE and SORTSIZE properties in **SASHome/SASFoundation/9.4/sasv9_local.cfg** be set to 120% of your largest data set. For example, if your largest input data set is 50 gigabytes, set MEMSIZE and SORTSIZE to 60 gigabytes.

On Windows deployments, you should add the `-MEMLIB` option to **SASHome\SASFoundation\9.4\nls\en\sasv9.cfg**.

For Linux deployments, you should put saswork on a RAM disk that is 90% of the size of the system memory or 230% of the size of the largest data set, whichever is bigger. For example, if your system has 128 gigabytes of memory, then your RAM disk should be 115 gigabytes. This value specifies the maximum size of the RAM disk. The space is not used until it is needed.

The following example commands provide a framework for how to create and assign the RAM disk. The commands must be completed in the order presented.

1. `sudo mkdir /fcmr_ramdisk`
2. `sudo mount -t tmpfs -o size=115g tmpfs /fcmr_ramdisk`
3. `sudo echo "tmpfs /fcmr_ramdisk tmpfs defaults 0 0" >> /etc/fstab`
4. `sudo mkdir /fcmr_ramdisk/saswork`
5. `sudo chmod 777 /fcmr_ramdisk/saswork`
6. `sudo ech "-WORK /fcmr_ramdisk/saswork" >> $sashome/SASFoundation/9.4/sasv9_local.cfg`

For systems where input data sources are very large or available memory is small, you might observe excessive memory paging to disk. In this case, it might not be preferable to use a RAM disk for saswork. Instead, you should add as many fast disk drives as possible to each grid node or the single-tier machine in a RAID5 or RAID6 LUN.

Disk Write operations to saswork create a bottleneck for projects with large data sources (typically greater than 10 gigabytes). If you observe disk Write operations headroom during big model training runs, you should consider increasing the number of pooled workspace servers from one to two or more. After increasing the number of pooled workspace servers, you should reevaluate whether model training run-time performance improved.

Stopped Models

When a SAS Factory Miner user stops a running model, it is stopped on the Job Execution Server. However, it continues running on the Pooled Workspace Server. When a model is running in MPP mode, a canceled model continues to consume resources on the MPP compute nodes.

Maximum Number of Threads

SAS Factory Miner enables you to specify the maximum number of threads that are run in any environment. However, when you use SAS Factory Miner in HPA or MPP mode with a Teradata grid, you cannot simultaneously create more than 15 models. The following steps ensure that SAS Factory Miner does not run more than 15 models, but you should set whatever value is appropriate for your installation.

1. Open SAS Management Console.
2. Expand the **Server Manager** plug-in and locate the application context that SAS Factory Miner is using. The installation instructions tell you to name this application context **SASAppFM**.
3. Expand the **SASAppFM** application context and the **SASAppFM — Logical Pooled Workspace Server** item.
4. Right-click **SASAppFM — Pooled Workspace Server** and select **Properties**.
5. On the **Options** tab, click **Advanced Options**.
6. On the **Load Balancing** tab, set the **Server process maximum** property to 15 or smaller.
Click **OK** to close the Advanced Options window.
7. Click **OK** to close the SASAppFM — Pooled Workspace Server Properties window.
8. Expand the **Object Spawner** plug-in and locate your SAS Factory Miner server.
9. Right-click your SAS Factory Miner server and select **Connect**.
10. On the **Options** tab, locate the **LogicalServer.SASAppFM — Logical Pooled Workspace Server.SASAppFM — Pooled Workspace Server.MaximumServers** property. This property should contain the old value, not the value that you specified in Step 6.
11. Right-click your SAS Factory Miner server under the **Object Spawner** plug-in and select **Refresh**. The **LogicalServer.SASAppFM — Logical Pooled Workspace Server.SASAppFM — Pooled Workspace Server.MaximumServers** property should now contain the value that you specified in Step 6.
12. Right-click your SAS Factory Miner server and select **Disconnect**.
13. Next, you need to update the Job Execution Services Application Configuration options. Select the **Application Management** ⇒ **Configuration Manager** ⇒ **SAS Application Infrastructure** ⇒ **Web Infra Platform Services 9.4** ⇒ **JobExecutionService** property. Right-click **JobExecutionService** and select **Properties**.
14. On the **Settings** tab, find the **Job Queue Maximum Threads** field. In this field, enter the same value that you specified in Step 6.

15. Also, on the **Settings** tab, click the **Pooled Workspace Server Properties** tab and locate the **Maximum Threads** property. Enter the same value that you specified in Step 6.
16. Click **OK** to close the JobExecutionService Properties window.
17. Close SAS Management Console.
18. Restart the SAS Web Application Server that is running SAS Factory Miner.

Exhausting the Pooled Workspace Server

Because SAS Factory Miner can create many concurrent jobs when creating models, users can exhaust all of the pooled workspace servers. This can drastically affect the performance of SAS Factory Miner and any other tasks that access the SAS server. To prevent or mitigate this issue, use the logical workspace server instead of the logical pooled workspace server for all operations other than creating models.

Complete the following steps to specify the logical workspace server:

1. Open SAS Management console and select **Application Management** ⇒ **Configuration Manager** ⇒ **Factory Miner Server Config** on the **Plug-ins** tab.
2. Right-click **Factory Miner Server Config** and select **Properties**.
3. On the **Settings** tab, find the **(Pooled) Workspace Server Name** property. If you followed the installation instructions provided in this book, this property should be set to **SASAppFM — Logical Pooled Workspace Server**. Instead of using the logical pooled workspace server, you want to use the logical workspace server.
4. Change the value of the **(Pooled) Workspace Server Name** property to **SASAppFM — Logical Workspace Server**.
5. Click **OK** to close the Factory Miner Server Config Properties window.
6. Restart the SAS Web Application Server that is running SAS Factory Miner. For example, this might be SASServer11 in an environment with multiple managed servers.

Define the SAS Factory Miner Data

Define a New Library

The steps in this section and the next section define the PURCHASES data set for use in SAS Factory Miner. This data set is used in the example that is found in the *SAS Factory Miner: User's Guide*. You can define additional libraries and register more data sets using the general steps presented in these sections.

Complete the following steps to define a data library for use in SAS Factory Miner:

1. Open SAS Management Console.
2. Expand the **Data Library Manager** to view the **Libraries** folder. Right-click **Libraries** and select **New Library**. The New Library Wizard window appears.
3. Select **SAS BASE Library**. Click **Next**.

4. In the **Name** field, enter a name for the library, such as **DMWebLib**. Enter a description that indicates this library is for use in SAS Factory Miner. Click **Next**.
5. Move the server that you defined for SAS Factory Miner from the **Available servers** field to the **Selected servers** field. If you followed the installation instructions in this book, that server is named SASAppFM. Click **Next**.
6. In the **Libref** field, enter **DMWebLib**, or whatever name you provided in step 4.
Click **New** to open the New Path Specification window. Click **Browse**. For Windows systems, navigate to **SASHome\SASFoundation\9.4\dmime**, select **sample**, and click **OK**. For Linux systems, navigate to **SASHome/SASFoundation/9.4/samples**, select **dmime**, and click **OK**.
Click **OK**.
7. Click **Advanced Options** to open the Advanced Options window. Enable the **Library is Pre-Assigned** option. Click **OK**.
Click **Next**.
8. Click **Finish**.

Register a Data Set

Complete the following steps to add a data source for SAS Factory Miner:

1. Right-click the **DMWebLib** library that you created in the previous section and select **Register Tables**.
2. Ensure that the proper **Libref**, **Engine**, and **Path** are assigned and select any options that apply for your data. Click **Next**.
3. Select the **Purchases** data set. Click **Next**.
4. Review your selection. Click **Finish**.

Your data is now available for use in SAS Factory Miner.

Note: Formats must be assigned when the workspace session starts. This is done either by modifying your autoexec file or adding start-up code in SAS Management Console.

Project Directory and File Information

Project Directory

Because the SAS Factory Miner project directory can become very large, it is recommended that you do not store projects in the configuration directory. You can change the project directory in SAS Management Console.

1. Open SAS Management Console.
2. Expand the **Application Management** plug-in.
3. Right-click **Factory Miner Server Config** and select **Properties**.
4. Click the **Settings** tab.

5. The **Projects Folder** property specifies where SAS Factory Miner projects are stored.

Migration Information

During migration, SAS Factory Miner project information is not migrated to the new location. This is true even when the project directory is located in the configuration directory.



Gain Greater Insight into Your SAS® Software with SAS Books.

Discover all that you need on your journey to knowledge and empowerment.

 support.sas.com/bookstore
for additional books and resources.


THE POWER TO KNOW.

SAS and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration. Other brand and product names are trademarks of their respective companies. © 2013 SAS Institute Inc. All rights reserved. S107969US.0613

