


# **SAS<sup>®</sup> Administration from the Ground Up**

**Running the SAS<sup>®</sup> 9 Platform  
in a Metadata Server Environment**



**Anja Fischer**

The correct bibliographic citation for this manual is as follows: *Fischer, Anja. 2019. SAS® Administration from the Ground Up: Running the SAS®9 Platform in a Metadata Server Environment*. Cary, NC: SAS Institute Inc.

**SAS® Administration from the Ground Up: Running the SAS®9 Platform in a Metadata Server Environment**

Copyright © 2019, SAS Institute Inc., Cary, NC, USA

978-1-64295-191-2 (Hardcover)  
978-1-63526-313-8 (Paperback)  
978-1-63526-729-7 (Web PDF)  
978-1-63526-727-3 (epub)  
978-1-63526-728-0 (mobi)

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, NC 27513-2414

March 2019

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.

SAS software may be provided with certain third-party software, including but not limited to open-source software, which is licensed under its applicable third-party software license agreement. For license information about third-party software distributed with SAS software, refer to <http://support.sas.com/thirdpartylicenses>.

# Contents

<b>About This Book .....</b>	<b>v</b>
<b>Acknowledgments .....</b>	<b>vii</b>
 <b>Chapter 1: Congratulations, You're the SAS Administrator! Now What?? .....</b>	 <b>1</b>
Why a starter-admin book? .....	1
What the book will and won't cover: .....	1
 <b>Chapter 2: Let the Admin Fun Begin: SAS 9.4 Architecture .....</b>	 <b>3</b>
Introduction .....	3
SAS Configuration Directories.....	5
SAS Client Tier.....	26
Data Tier .....	26
Summary.....	26
 <b>Chapter 3: Administration Tools .....</b>	 <b>27</b>
Introduction .....	27
The SAS Jedi: SAS Management Console .....	28
SAS Environment Manager.....	40
SAS Deployment Manager.....	58
SAS Web Administration Console.....	73
Summary.....	75
 <b>Chapter 4: Do I Know You? A Quick Review on Users, Groups, and Roles in SAS 9.4 .....</b>	 <b>77</b>
Introduction to User, Group and Role Management.....	77
Users.....	80
Summary.....	90
 <b>Chapter 5: Metadata Library Administration in SAS 9.4 .....</b>	 <b>93</b>
Introduction .....	93
Libraries in Base SAS .....	93
Creating Libraries in SAS Metadata .....	95
Summary.....	106
 <b>Chapter 6: SAS 9.4 Metadata Security.....</b>	 <b>107</b>
Introduction .....	107
Authorization .....	109
Best Practice: The Golden Rules for Metadata Security Model Design.....	120
Security Testing Options .....	133
Summary.....	136

## **Chapter 7: Backup and Restore: Should Have, Could Have, Would Have ... Famous Last Words ..... 137**

Introduction .....	137
Backup, Backup and Backup Again .....	138
Backup and Restore Tools .....	140
Summary.....	153

## **Chapter 8: SAS Administration Tasks and Housekeeping..... 155**

This is it – Finish Line in Sight: SAS Administration Tasks and Housekeeping.....	156
Your List of SAS 9.4 Administration Tasks for Starters .....	156
Now What?.....	171
Summary.....	172

# About This Book

---

## What Does This Book Cover?

This book is about the basic principles of SAS 9.4 platform administration. It is a starter guide for new SAS administrators, helping you to turn into a happy, calm and confident SAS admin. The book provides you with a light entrance into the world of SAS administration, without wading through the documentation. This book does not cover SAS Viya administration.

---

## Is This Book for You?

If you are new to the SAS 9.4 admin job or are an advanced admin who wants to make sure you know all the admin basics and admin tricks, this book is for you.

---

## What Should You Know about the Best Practices?

This book provides recommendations and best practices around the most important SAS 9.4 platform administration topics a SAS Platform administrator should be familiar with.

---

## Additional Resources

Please find an appendix full of handy links and references on the author's page at <http://support.sas.com/fischer>.

---

## We Want to Hear from You

Do you have questions about a SAS Press book that you are reading? Contact us at [saspress@sas.com](mailto:saspress@sas.com).

SAS Press books are written *by* SAS Users *for* SAS Users. Please visit [sas.com/books](http://sas.com/books) to sign up to request information on how to become a SAS Press author.

We welcome your participation in the development of new books and your feedback on SAS Press books that you are using. Please visit [sas.com/books](http://sas.com/books) to sign up to review a book

Learn about new books and exclusive discounts. Sign up for our new books mailing list today at <https://support.sas.com/en/books/subscribe-books.html>.

Learn more about this author by visiting her author page at <http://support.sas.com/fischer>. There you can download free book excerpts, access the appendix, read the latest reviews, get updates, and more.



# Chapter 2: Let the Admin Fun Begin: SAS

## 9.4 Architecture

<b>Introduction .....</b>	<b>3</b>
<b>SAS Configuration Directories .....</b>	<b>5</b>
SASHome .....	5
Lev1/Levn .....	7
SAS Tiers: The three plus one SAS Tiers in a metadata-managed environment .....	9
SAS Server Tier .....	12
About the Workspace Server, Stored Process Server and all other SAS Servers .....	15
SAS Connection Profiles .....	19
SAS Web Application Server and SAS Web Server (http server) .....	24
Cache Locator.....	25
JMS Broker.....	25
SAS Environment Manager .....	25
Java Runtime Environment.....	25
<b>SAS Client Tier .....</b>	<b>26</b>
<b>Data Tier.....</b>	<b>26</b>
<b>Summary .....</b>	<b>26</b>

---

### Introduction

The starting point for SAS administration is the architecture: what are the components of a *SAS deployment*, how does it look? With a good understanding of the SAS architecture, you'll be able to tackle the responsibilities and tasks that come with SAS administration. Understanding the architecture means you know where to look if you need to make any changes, troubleshoot, and the like.

The best way to explain the SAS 9.4 environment architecture is an analogy to the architecture of a house. Envision the following:

You buy a piece of land (infrastructure/hardware). You want to build a new house (software) on your land. How do you build your house (install your software)? You need an architect (an admin). Sometimes there is one architect, sometimes there are more than one (one admin vs many admins). Once you decide on the shape, storage, levels, house color etc. first thing the home builders do is lay down a foundation (SAS metadata server). On top of that, the walls are built for either a one-story or multiple-story house (distributed or non-distributed SAS environment).

Once the house is built and is move-in ready, you have bedrooms, guest room, kid's room, office, kitchen, bathroom, etc., each of which has a different purpose (a different "task"). The



#### 4 *SAS Administration from the Ground Up*

different floors, the different rooms, are your SAS servers, each of which fulfills different tasks, different needs - a different purpose.

Last but not least, think about the different objects in these rooms: toys, towels, beds, plates, glasses, toothbrushes etc. All these objects can be compared with data sources: SAS data sets, DBMS, Hadoop etc. Knowing what is in each room helps find what you are looking for. Same with SAS architecture: if you know and understand the architecture, you know where to look.

Now, your house is completed, and you have moved in. You're done. Right?

Very wrong. After you move in to this great new home, you must maintain it to keep it great, clean and beautiful. Think about it: You must wash the windows, change the air filters, check the air conditioner at least twice a year, do dishes, vacuum, etc.— some of the tasks more frequently than others. Same concept with SAS: once it is installed, you must maintain it to keep it clean, healthy and good looking. So, let's apply the house analogy to SAS.

Let's start with the different install flavors (single house versus townhouse, single story versus multiple stories). SAS can be installed either as a SAS Foundation install or as a metadata managed install.

**Note:** In this book, we will focus on metadata managed deployments only!

SAS Foundation is your basic install, think Base SAS. A metadata managed install is the SAS 9 Intelligence Platform, with much more than Base SAS. You might have SAS Visual Analytics, Grid, SAS solutions, SAS Add-In for Microsoft Office, etc. With SAS Foundation, your users work on their personal machines, or use Remote Desktop or Citrix. A SAS Foundation install does not involve a centrally metadata managed system. In a metadata managed install, your users work on the dedicated SAS server.

The two different SAS deployments can be installed on physical or virtual machines.

**Note:** For every SAS solution, every Grid install, SAS Visual Analytics (and more), the SAS Platform administration applies. Even though Grid, VA and SAS solutions have additional, product-specific administration tasks: EVERY PRODUCT IS BASED ON SAS 9.4 PLATFORM ADMINISTRATION!

Let's take a peek at the SAS Configuration. We will only cover it briefly, to give you some very basics.

---

## SAS Configuration Directories

After SAS is installed, you'll find two different SAS directories. One called **SASHOME** and another one, called **Lev1** (aka configuration directory) for the metadata managed, site-specific configuration for your SAS 9 platform services.

New admins are often taken by surprise that SAS has two directories, and don't quite know what to do with them. I totally get it. So, let's see whether we can shed some light on this "dual" directory situation.

---

### SASHome

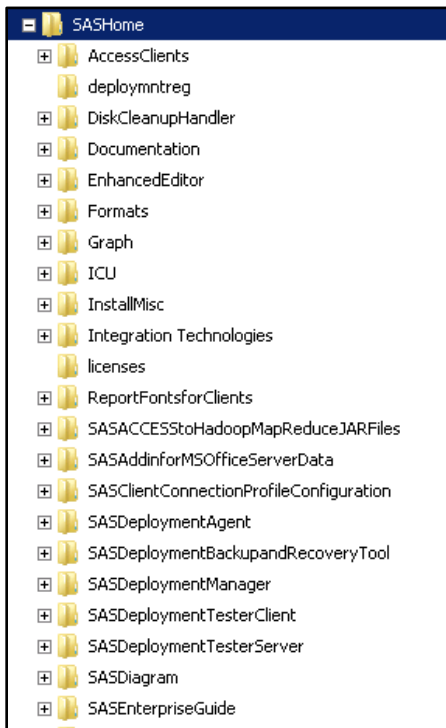
SASHOME includes subfolders for all your SAS desktop clients and SAS web clients. This directory (aka SAS root folder) is located per default at:

- C:\Program Files\SASHome on Windows, and
- /usr/local/SASHome on Linux/Unix.

Depending on the way SAS is installed, !SASHOME can be at another location

On Windows, for example, SASHOME might look like:

**Figure 2.1: SASHOME directory**



**Note:** Depending on the products you have licensed, you might have different folders.

Aside from executables and configuration files in each respective client folder, you find some other cool things, such as examples, SAS programs and data that is specific to that client.

Take SAS Enterprise Guide for example. Look at

**SASHome\SASEnterpriseGuide\7.1\Sample** and you'll find code examples, data sets, and Enterprise Guide example projects.

If users run tests and do not want to touch production, or new users must come up to speed with SAS Enterprise Guide, these are some examples of situations where this test data might come in handy.

Another example is the *SASFoundation* folder, in which you can find example data sets, programs, catalogs and views: `\SASHome\SASFoundation\9.4\core\sample`

**Note:** The SASFoundation folder also stores your setinit information:  
**!SASHome\SASFoundation\9.4\core\sasinst**

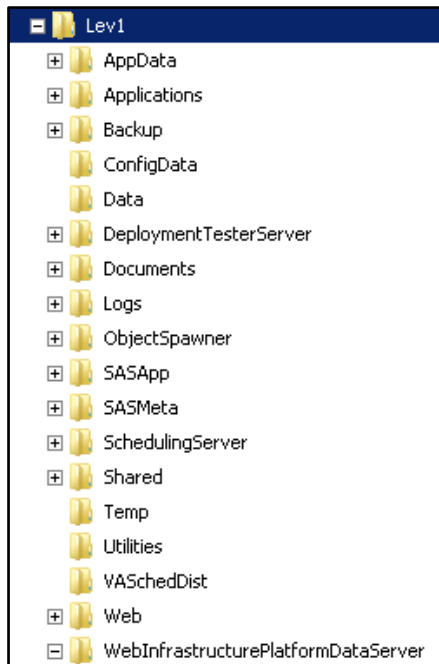
The setinit is your license file. This file includes the products your company has licensed and the date when the license expires.

## Lev1/Levn

Lev1/Levn is the metadata managed, site-specific configuration for your SAS 9 platform services. The default path for the configuration directory is `\SAS-config-dir\Lev1\` where “SAS-config-dir” is the path that you chose during the deployment.

Just as with SASHOME, the configuration directory `\SAS-config_dir\Lev1\` includes configuration files, scripts, etc. An example for a site-specific configuration directory is shown in Figure 2.2.

**Figure 2.2: Lev1 Directory**



**Note:** Depending on the products you have licensed, the directory structure might look differently.

**Note:** In some deployments you might find a *Lev2* and/or *Lev3*. In such cases it simply means that one might have set up a dev, test, and prod environment, or, runs different SAS 9 versions in parallel.

The following lists the content of the *Levn* subdirectory.

### Contents of the *Levn* Subdirectory

(Resource: SAS® 9.4 Intelligence Platform: Administration / System Administration, available at:

<https://go.documentation.sas.com/?cdcId=bicdc&cdcVersion=9.4&docsetId=bisag&docsetTarget=p1oa9ysgpowj4vn19o67gc7xrrr0.htm&locale=en>)

Subdirectory or File	Description
<b>AppData</b>	Contains indexes and the repository configuration file for the SAS Content Server, and data that is installed for the use of specific applications (for example, SAS BI Dashboard).
<b>Backup/Vault</b>	Is the default location for backups that are created by the Deployment Backup and Recovery Tool.
<b>ConnectSpawner</b>	Contains the management script, configuration files, and logs for the SAS/CONNECT spawner.
<b>Data</b>	Can be used to store user data.
<b>DeploymentTesterServer</b>	Contains files that are used by the Deployment Tester plug-in to SAS Management Console.
<b>Documents</b>	<p>Contains <code>Instructions.html</code>, which contains post-installation configuration instructions; <code>DeploymentSummary.html</code>; <code>ConfigurationErrors.html</code>; and other application-specific documents</p> <p>.Tip:</p> <p>The <b>instructions.html</b> file includes all information about how SAS was installed, errors or warnings that might have occurred during the install, links, ports and other helpful information. You can look at that file to find out about configuration paths, log file locations, links etc. It can be helpful to familiarize yourself with your SAS install. Aside from the <code>instructions.html</code> file, there are other helpful information available. The next table lists some of them as examples.</p>

Subdirectory or File	Description
<b>Logs</b>	Can be used as a common directory for server and spawner logs, if you selected this option during a custom installation. By default, each server has its own separate log directory.
<b>Logs/Configure</b>	Contains logs that are created by the SAS Deployment Wizard.
<b>ObjectSpawner</b>	Contains a management script, configuration files, and logs for the object spawner.
<b>SASApp</b>	Contains management scripts, configuration files, and logs for SAS Application Server components
<b>SASMeta</b>	Contains management scripts, configuration files, metadata repositories, logs, and other files and directories for the SAS Metadata Server. .
<b>ShareServer</b>	Contains the management script, configuration information, and log files for the SAS/SHARE server.
<b>Utilities</b>	Contains XML files that are used as input to the SAS Deployment Wizard and utilities that are associated with this configuration instance.
<b>Web</b>	See Contents of the Web Subdirectory.
<b>WebInfrastructurePlatformDataServer</b>	Contains the management script, configuration information, and log files for the SAS Web Infrastructure Platform Data Server.
<code>generate_boot_scripts.sh</code>	Is a script that is used to regenerate the <code>sas.servers</code> script on UNIX.
<code>metadataConfig.xml</code>	Provides information for application server components to use when they connect to a clustered metadata server.
<code>sas.servers</code>	Is a management script that is used on UNIX to start, stop, or restart all servers on the machine in the correct order, or to display the status of all servers on the machine.
<code>sasv9_meta.cfg</code>	Specifies metadata server connection information for the SAS OLAP Server and SAS/SHARE server.

## SAS Tiers: The three plus one SAS Tiers in a metadata-managed environment

The SAS *Levn* configuration is only one part of the SAS Platform. It is also important to know that your environment is an n-tier architecture, which means that all components that

make up your SAS environment, can be distributed over multiple computer resources. Each tier component, each tier part, performs only the work it is responsible for.

Depending on the number of servers you have available, the SAS tiers can be installed across multiple machines or on a single machine. Going back to our analogy at the beginning: each room fulfills a different purpose, “stores” different objects, serves different members of a family. All of the rooms together make your house (SAS deployment). Now, what are these tiers?

Getting to grips with the SAS tiers truly is a fundamental requirement for a good SAS administrator. In SAS 9, we have three tiers, plus one. Plus one, because the fourth tier is not a SAS tier, but is an important element. Let’s investigate each tier a bit further.

### **SAS Servers Tier aka Compute Tier**

Simply put: SAS servers perform SAS processing on your data. They are SAS Server processes running on one or more physical or virtual server machines.

Important – and often misunderstood: The tiers are not actual physical machines, but processes, where each process (SAS server) has its own tasks. One function is providing your users with the data they are requesting, or, running a Stored Process, or creating a report. And that’s really all there is to it: the tier these SAS processes (SAS servers) are running on is referred to as Server Tier or Compute Tier.

Next, we have the web tier, called the middle tier.

### **Middle Tier aka Web Tier**

The middle tier – also called the web tier – coordinates the web traffic. It enables access to data and functionality using web clients, a browser. Think SAS Studio (users using a web browser), SAS Environment Manager (web client) or any other web-based SAS APIs. We will revisit the middle tier in a little bit.

### **Client Tier**

The client tier runs your desktop clients and web browsers. Examples for SAS clients are SAS Enterprise Guide, SAS Add-In for Microsoft Office, SAS Enterprise Miner, and so on.

### **Data Tier (Data Sources) – The “Plus One” Tier**

The data tier is where the data sources are stored. It is different from the SAS tiers we just described because even though it is important to SAS, it does not come from SAS. It is not used for SAS to run but for users to consume using SAS.

A data tier can be any machine that stores data that you want to access from within SAS.

Data sources can be SAS data sets, OLAP cubes, web content, DBMS data (SAS can access databases such as Oracle as one example out of many), and more. We will talk more about data in Chapter 5, Metadata Library Administration.

Table 2.1 depicts the three plus one tiers. Even though the tiers are pictured in four different boxes, the boxes do not represent physical machines. It simply pictures the different layers a

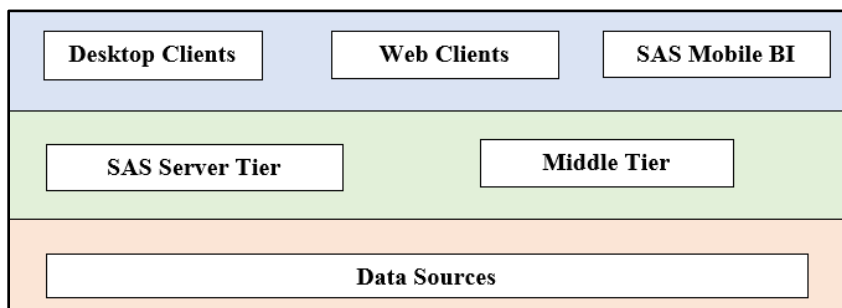
SAS metadata deployment consists of, whether it is installed on one machine or multiple machines.

**Table 2.1: Architecture of the SAS Intelligence Platform**

Data Tier	SAS Server Tier	Middle Tier	SAS Client Tier
SAS Data Sets	SAS Metadata Server	SAS Web Server	Desktop Clients: SAS Add-In for Microsoft Office
OLAP Cubes	SAS Workspace Server	SAS Web Application Server	SAS Data Integration Studio
SAS Scalable Performance Data Server (SPDS)	SAS Pooled Workspace Server	SAS Web Infrastructure Platform	SAS Enterprise Miner
SAS Web Infrastructure Platform Data Server	SAS OLAP Server	<b>Task:</b> providing services and applications for SAS web applications. <u>It includes:</u>	SAS Forecast Studio
Third-Party Data Sources (such as Oracle, Teradata, etc.)	SAS Stored Process Server	SAS Content Server to store digital content (such as reports) and	SAS Enterprise Guide
Hadoop	<i>These servers are running SAS processes for distributed clients</i>	other infrastructure applications and service ( <i>such as SAS Deployment Backup and Recovery tool, and more</i> )	SAS Information Map Studio
Enterprise Resource Planning (ERP) Systems		Web Clients: (run in an instance of the SAS Web Application Server). The SAS web clients are:	SAS Management Console
		SAS Web Report Studio	SAS Model Manager
		SAS Information Delivery Portal	SAS OLAP Cube Studio
		SAS BI Portlets	SAS Workflow Studio
		SAS BI Dashboard	JMP
		SAS Help Viewer for the Web	Other SAS analytics products and solution
		Other SAS Web Applications and Solutions	Web Browser to surface web Applications
		SAS Environment Manager (server process that includes a web application server, providing a web-based administrative interface)	Mobile Devices, if applicable, to view certain type of reports

Figure 2.3 shows this from a very simplified layer perspective.



**Figure 2.3: A layered view of SAS platform architecture.**

To take a closer look at these tiers, take a look at the Architecture Overview section in *SAS® 9.4 Intelligence Platform: Administration / SAS Intelligence Platform: Overview*, available at: <https://go.documentation.sas.com/?cdcId=bicdc&cdcVersion=9.4&docsetId=biouv&docsetTa rget=titlepage.htm&locale=en>

---

## SAS Server Tier

Let's come back and talk a bit more about the SAS Server tier. I would like to spend some time discussing its components because as a SAS administrator, this is super important!

To refresh your memory, SAS servers are SAS Server processes running on one or more physical or virtual server machines. The SAS server tier is nothing but a bucket of SAS processes that are based on "sas.exe". Every "sas.exe" has a different role or functionality, such as:

- metadata server (in-memory)
- workspace server (interactive sessions)
- stored process and pooled workspace server (trusted sessions)
- if you have SAS Visual Analytics for example, the SAS LASR Analytic Servers (in-memory for SAS Visual Analytics). We will not address SAS Visual Analytics; this is simply meant as an example.

Let's discuss each of these in turn.

## Metadata Server

The metadata server is the heart of your environment, the foundation. If it's not working, SAS is not working. Going back to the house-building analogy: if the foundation is weak or breaks, the house will crumble.

The metadata server is an in-memory server. It keeps your environment running and stores all your assets: metadata about where to find your assets, where assets can be data, users and groups, etc. It is a centralized resource for storing, managing, and also providing metadata for all your SAS applications, for all your users.

In-memory: when your users request data, a copy of the physical data is stored into memory. From there on, your users are going “against” memory. This speeds up the process so that the access time is quicker. Memory is flushed when you pause and resume, or, stop and restart the metadata server.

When speaking about metadata assets, we are referring to libraries, users, groups, SAS folders – everything you create using SAS Management Console or SAS® 9.4 Open Metadata

Interface (metadata server programming language). SAS applications also create and manage metadata.

To learn more about the SAS 9.4 Open Metadata Interface, take a look at <https://go.documentation.sas.com/api/docsets/omaref/9.4/content/omaref.pdf>

These assets are stored in a *metadata repository*. A repository is like a box in which you save all your belongings. Your main repository is *Foundation*. You can also create custom and project repositories. When SAS is installed, only one repository is created, which is the Foundation one.


A custom repository can be used – as one example – in cases where you might have very sensitive data that should only be available to a certain group of users. The custom repository and its contents could be created in a directory where only these users have access to.

Project repositories are used for Change Management and are available for SAS Data Integration Studio only. Users can check out and lock metadata so that metadata can be modified and tested, to be checked back in afterward, which then unlocks the metadata.

In this book, we will focus on the Foundation repository. To learn more about custom and project repositories, check out About Repositories [at https://go.documentation.sas.com/?cdcId=bicdc&cdcVersion=9.4&docsetId=bisag&docsetTarget=p1b7gxgkbu04zbn1dnhh3pb8c6zs.htm&locale=en](https://go.documentation.sas.com/?cdcId=bicdc&cdcVersion=9.4&docsetId=bisag&docsetTarget=p1b7gxgkbu04zbn1dnhh3pb8c6zs.htm&locale=en)

When you log on to SAS Management Console, *Foundation* is chosen per default as shown in Figure 2.4.

**Figure 2.4: SAS Management Console**



**Note:** The path defaults to **SAS-conf-dir/Lev1/SASMeta/MetadataServer/MetadataRepositories/**. Do not rename or delete the metadata server repository path and never move, delete, or modify data sets in the MetadataRepositories and rposmgr directories.

You can remove metadata objects via SAS Management Console or by using metadata programming features using the SAS® 9.4 Open Metadata Interface. Before you delete or remove any metadata objects, I recommend the following information from the sasCommunity.org Planet: “*Where is my stuff? Documenting what is stored in SAS Metadata*” available at: <http://www.sascommunity.org/planet/blog/category/metadata/>.

Note: if you are using the printed book version, I realize that having to type a link into a web browser is a bit laborious, but this is one of the articles that are really helpful, and hard to find if you don’t dig deep. Typing it is worth it.

## **Metadata server and system access**

You can use the metadata server’s authorization model to control access to your SAS assets (aka metadata objects). SAS security does not include protection of configuration files or any non-SAS-related content. In the Appendix of the book, you will find an overview of the operating system protections for Windows and Unix/Linux. Chapter 6, SAS 9.4 Metadata Security will cover metadata server authorization.

## **How about a clustered metadata server? Let’s talk high availability!**

You have the option to set up the SAS Metadata Server as a clustered configuration. A metadata server cluster provides redundancy and high availability, which helps you make sure your environment will continue to operate should one metadata server go down.

You can implement metadata server clustering at any time. It does not have to be set up during the initial install.

A common question is if the SAS clients know when one cluster node goes down and another one picks up. There is no configuration for your SAS clients necessary. SAS clients, such as SAS Management Console, keep a list of the metadata server cluster nodes, which is updated each time a client connects to the cluster.

Where changes are required is for the SAS Application Server tier. Application servers such as the Object Spawner, workspace server, stored process server, etc. don’t understand that there is a cluster all of a sudden. Application servers use a configuration file – called *metadataConfig.xml* – which tells them about the metadata server. So as long as you don’t make changes to this file, the SAS application servers still assume that there is only one metadata server. The SAS middle tier applications keep a list that is automatically updated when the Web Infrastructure Platform web application starts. This is all described in the information for metadata server clustering, which you will find in the resource references in the Appendix of this chapter.

In addition to the official SAS documentation about metadata server clustering, you will find some other great resources such as papers and blogs.

## How do clients interact with the metadata server?

When users start a SAS client, such as SAS Enterprise Guide, a connection profile is used to connect to the metadata server. I will talk about Connection Profiles a little later in this chapter.

## Metadata Server logging

The default location for the Metadata Server log file is:

```
SAS-config-dir\\Lev1\SASMeta\MetadataServer\Logs
```

SAS 9.4 uses a standard logging facility to perform logging for the metadata server and all other SAS servers. Generally, the default logging information are sufficient, as it provides you with information, warnings and errors. SAS Technical Support might ask you to increase the logging level if more precise information is needed for troubleshooting. You can certainly modify the logging levels at any time. Just keep in mind that increasing the information the metadata server writes to the log, the metadata server log file size will grow. Maintaining log files will be discussed further in the housekeeping chapter.

**Note:** If the reason for considering enabling additional logging is because you want more information for auditing and monitoring, logging will not fulfill your needs. In Chapter 3, SAS Administration Tools, we will talk about SAS Environment Manager, which provides you with some great options for monitoring or auditing.

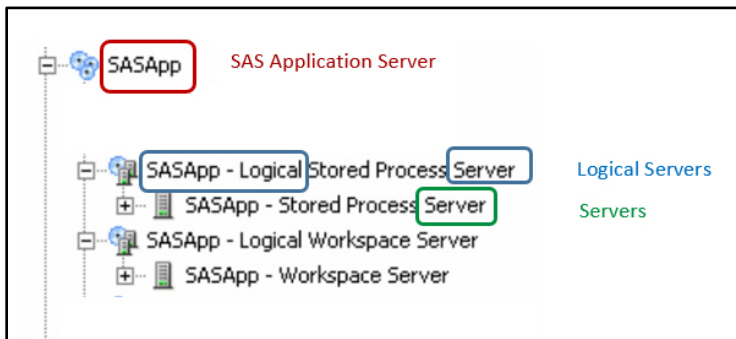
**Important:** If you look at the individual log folders for your SAS servers, you will notice that the workspace server does not have any log files. A workspace server log is written for each individual user. Let your users work for a day with a workspace server log enabled and you'll end up with an abundance of log files. Workspace server logs are usually only enabled when SAS Technical Support needs information for further troubleshooting. Enabling workspace server logging just out of curiosity can affect the performance.

---

## About the Workspace Server, Stored Process Server and all other SAS Servers

These servers are called SAS Application Servers. The SAS Application Servers are a set of metadata objects using application server specific configuration files that are automatically created when SAS is installed. You can look at the SAS Application Server definitions in SAS Management Console, Server Manager. The SAS Application Server Context is called SASApp by default.

Most SAS Application Servers are structured with a Logical Server component and a Server component. Figure 2.5 is an example for the Stored Process Server.

**Figure 2.5: Stored Process Server Structure**

The SAS Application Servers are simply SAS processes, fulfilling all different types of user requests. The workspace server is used for general client user interactions. The stored process server is used for stored processes, and so on.


**Note:** You must not rename the name *SASApp*. Configuration files use this name reference and its configuration settings for SAS client interactions. Renaming *SASApp* could result in your users not being able to run jobs within SAS clients.

Underneath the *SASApp*, you can see the various dependent SAS Application Server objects, and the way they are structured. Almost all server components consist of a *Logical Server* component and a *Server* component.

## Logical Servers and Servers

SAS servers, such as the workspace server or stored process server for example, can be grouped into logical objects. That means, you can have multiple servers grouped together under one logical application server.

Let's stick with the SAS workspace server for a moment. An example for server definitions under a **Logical Server** definition is:



But it could also look like



As you can see in this example, a *Logical Server* can have one Workspace Server definition or multiple. In this case, I created two additional workspace servers based on the department I

work in, called Customer Success. In this department, we have a group of Systems Engineers and Customer Success Managers. We separate the workspace servers for both groups for purposes such as:

- You want a certain group of users to run their jobs in a specific application server. Every time a user (such as a Systems Engineer) runs a job, the job is executed by a specific workspace server (*Systems Engineers Workspace Server*).
- You can also use it if you want to assign certain content. We could assign specific SAS libraries to each individual group, as in our case, libraries for Systems Engineers and libraries for Customer Success Managers.

Additional logical servers are created using the SAS Deployment Wizard.

Another example for a scenario in which the creation of additional servers might make sense is in a dev, test and prod environment.

Creating dev, test, prod environments (additional SAS 9.4 instances) is a big topic in itself, and the example for creating additional servers for a dev, test, prod environment is just one example. Another way to set up multiple SAS 9.4 instances is with the SAS Migration Utility. The Appendix of this chapter will provide you with the resource reference, talking about how to copy existing deployments to create additional SAS 9 instance using the SAS Migration Utility.

There are many great discussions about this topic on the SAS Administration and Deployment Community. One of which I would like to highlight is *Running two SAS Environments on the same Linux machine*.

The question was if there is a best practice for deploying SAS 9.4 on the same machine to create additional separate environments. In this case, the user was looking to create a second SAS 9.4 instance on the same machine. Paul Homes, an expert and owner of the SAS partner metacoda, provided some great information that I would like to share with you. The question was for Linux, but the concept can be applied to other operating systems flavors as well. I would like you to focus on the pros and cons described. Paul's answer to the question:

*You can definitely run 2 environments on the same machine (given the machine is suitably sized in terms of memory, CPU, storage, etc.). Use the SAS Deployment Manager to do an install and deployment of Lev1 and then run it again to do another deployment of Lev2. You can choose different Lev numbers as long as they are different for each environment on the same machine. By choosing a different Lev number all of the default port numbers will be automatically chosen to avoid conflict.*

*Of course, there is also the option to have multiple virtual machines on the same physical machine.*

*If you are not installing new 'empty' environments then the SAS Migration Utility can be used to create a package based on an existing SAS environment so you can use it during the deployment of a new environment. Alternatively, you can use the export/import tools to migrate metadata and content using SAS package files post-install.*

## 18 SAS Administration from the Ground Up

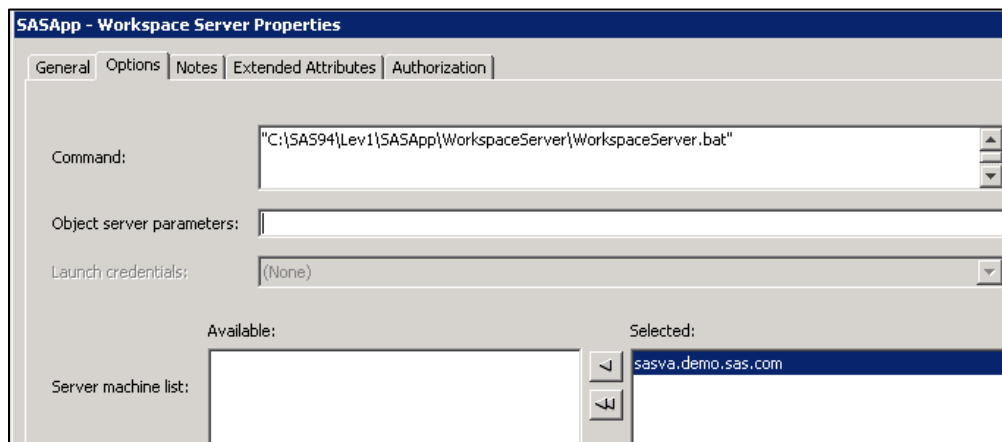
*A benefit of having both environments on the one machine is one less machine to procure and manage. A downside is that the machine will need to be bigger to support both environments. Being less isolated can also make it harder to do changes to one environment without potentially impacting the other environment. (...) Imagine a DEV and PROD on the same machine. You wouldn't be able to test opsys and SAS patches on DEV prior to their application to PROD. You wouldn't be able to reboot the DEV machine without rebooting the (same) PROD machine.*

*As to whether the pros outweigh the cons or vice-versa depends on how you plan to use the environments. I would definitely recommend consulting SAS Professional Services or a local SAS Partner if you want some help with planning.*

I went a bit of track there, let's get back to our SAS Application Servers ...

SAS Application Servers are accessed either by desktop clients or, by web applications that run in the middle tier. Logical Servers are “holding” the Servers that you have per default, plus, any custom servers should you decide to add them. The workspace server is just an example. It is the same concept for other SAS Application Servers as well.

Logical Servers contain one – and one only – application server definition, such as the workspace server. The object name SASApp – Workspace Server or SASApp – Stored Process Server etc., hold the information for the server that executes SAS code. If you look at the properties for the SASApp – Workspace Server for example, you'll see the following:



The file WorkspaceServer.bat includes the information needed to start the workspace server and the selected server machine is the machine where the server runs on. Customers create additional servers to separate group access, as one example.

The hardcopy readers amongst us are probably getting close to hit me with my admin book, because I would like – once more – highlight a paper that goes with the SAS workspace server and might come in handy should you ever experience any delays with the SAS workspace server initialization. But remember all these links can be found on my author page in a handy pdf. The paper is called: Tracking Down the Culprit of a SAS® Workspace Server Initialization Delay, available at <https://www.sas.com/content/dam/SAS/support/en/sas-global-forum-proceedings/2018/2003-2018.pdf>.

---

## SAS Connection Profiles

While we are talking about the SAS servers and services, users, and client products such as SAS Enterprise Guide, I believe that SAS connection profiles are worth mentioning.

### Logging In

When users log on to a client, or an admin logs on an admin tool – whether it is a desktop or web client – a connection profile is used. Then, users or admins connect to the metadata server. In order for the web or desktop clients to “find” the metadata server, we need the metadata server “address,” aka connection profile. The connection profile simply stores user credentials, ports, and server names. You can find the profiles in the following location for SAS Data Integration Studio, SAS Information Map Studio, SAS Management Console and SAS OLAP Cube Studio:

- Windows Vista or later: C:\Users\user-name\AppData\Roaming\SAS\MetadataServerProfiles
- UNIX: /user-home/SASAppData/MetadataServerProfiles.

The client profiles have the extension **swa**. Here is a snippet of a client profile:

```
port=8561
```

```
userid=sasdemo
```

```
Name=SASDemo
```

```
password={sas002}1D57933958C580064BD3DCA81A33DFB2
```

```
host=machine_name
```

In this profile, the user name and password are stored, which means the user will not be prompted. The users have the option to check a box during login that saves the user ID and password in their profile.



**Tip:** If you would like to avoid that users have the option to check this check box, you can do the following:

On your metadata server machine, go to

`sas_config_dir\Lev1\SASMeta\MetadataServer` and open the file `omaconfig.xml`.

Change the value for `SASEC_LOCAL_PW_SAVE` from 1 to 0, where 1 is YES and 0 is NO.

Save your changes and close the file.

Restart your metadata server for the changes to take effect. Please keep in mind that the restart of the metadata server will throw out all your users, meaning, their work will be interrupted. For that reason, you might want to choose a time where there is the least user traffic.

This will disable the check box to save user ID and password from the profile.

## Quick excursion to SAS encryption

Looking at the client's .swa file, you might notice the password:

```
password={sas002}1D57933958C580064BD3DCA81A33DFB2.
```

SAS encrypts password at rest and in transit. There are several encryption mechanisms available in SAS. Here, you see `sas002`, which is the default SAS encryption called *SASProprietary*, which is a fixed coding algorithm with medium security.

OK, that was the SAS Application Servers. Next, I would like to take a moment and look at the SAS Object Spawner.

## Object Spawner

Another important SAS component that we must talk about when talking about SAS application servers is the Object Spawner. An object spawner runs on each machine where you want to run a workspace server, pooled workspace server or stored process server.


The Object Spawner's task is it to launch a workspace server, pooled workspace server or stored process server, whenever one is requested. If a user accesses a table in SAS Enterprise Guide to work with it, the workspace server is used to execute the user's job, right? Not quite. The component that actually starts a workspace server session is the Object Spawner.

Before the Object Spawner starts any of these application servers, it establishes a communication with the metadata server to check whether the requesting user has a valid user ID.

To be able to have a communication between the Object Spawner and the metadata server, the object spawner uses a configuration file that includes the information needed to access the metadata server. The configuration file is called **metadataConfig.xml** and is located at **SAS-config-dir\Lev1\ObjectSpawner**.

It includes the metadata server machine, the metadata server port and other information. Think of it as if you are giving someone your address to find you.

The SAS documentation uses the following figure to show how the spawner obtains metadata:



❶ The object spawner accesses the metadataConfig.xml

The object spawner connects to the SAS Metadata Server for configuration information.

It is like me asking my German pal where I can get the best Bratwurst. She gives me the name of a store and I take that information to get to the store. (quite the example, isn't).

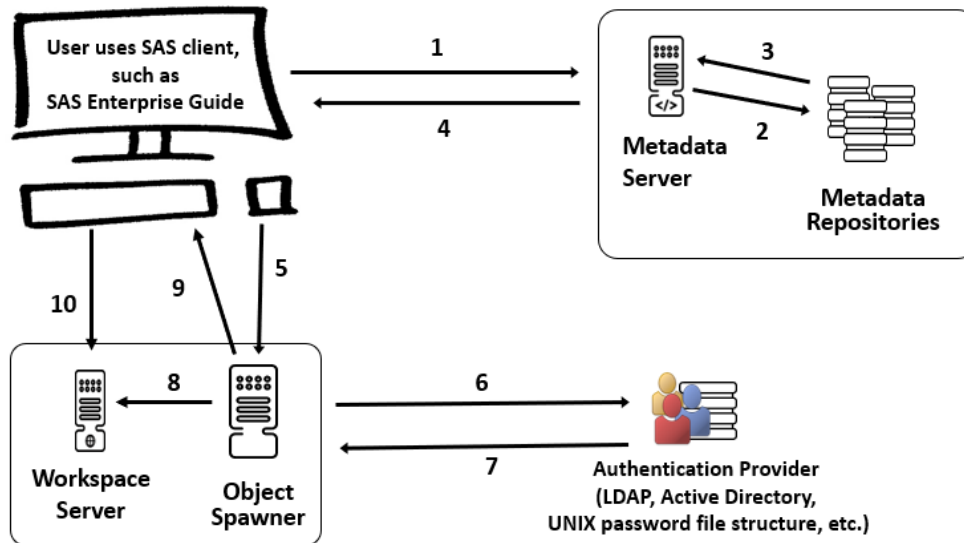
## Authentication and starting SAS Servers

### Example for Workspace server:

Before the Object spawner starts a workspace server, the metadata server provides credential information and details about how to start the server. In a default scenario, the object spawner uses host authentication, which means the workspace server is started under the users' credentials.

After the client – such as SAS Enterprise Guide – is exited, the workspace server session is closed. You can use token-authentication instead of host authentication. Token authentication uses a shared user and generates a single-use identity token.

Let's take a look at it:



The steps in detail:

1. Using the already established connection to the metadata server, the SAS client, here the user is using SAS Enterprise Guide, requests access to a workspace serve (1)
2. In step 2, the metadata server searches the metadata (metadata repositories) for the workspace server in question.
3. The metadata server then gets the machine name hosting the workspace server, the Object Spawners's port and an authentication domain that is associated with the workspace serve (we will address authentication domains in the security chapter).
4. This connection information is returned to SAS Enterprise Guide.
5. SAS Enterprise Guide uses the connection information to make the request for a workspace server. If the authentication domain for the server matches that of the initial inbound login, SAS Enterprise Guide passes along the credentials as well.
6. The authentication domain comes into play for the credentials that are being used. If the user ID is not associated with the default authentication domain (called DefaultAuth), SAS Enterprise Guide searches its in-memory list of credentials to find the user's credentials with the appropriate authentication domain. If no user credential is found, SAS Enterprise Guide queries the metadata server for credentials for this user for that particular authentication domain (outbound login). If none is found, the user will be prompted for credentials.
7. The object spawner sends the user's credentials to its authentication provider. The default authentication provider is the host. As shown in the graph above, different authentication providers can be used.
8. The authentication provider verifies that the credentials are valid.

9. The object spawner launches the workspace server. It uses the launch command that was retrieved from the metadata at start-up. The workspace server runs under the credentials provided by SAS Enterprise Guide, that have been authenticated by the host.
10. The object spawner provides SAS Enterprise Guide with a TCP connection to the workspace server session.

SAS Enterprise Guide communicates directly with the workspace server.

Now, the user submits requests with SAS Enterprise Guide for processing. Results are returned to SAS Enterprise Guide. Easy as that: Code submitted, results returned

My Bratwurst example sounded simpler.

There are great troubleshooting tips available for SAS servers. Because I am afraid to add yet another note apologizing to the hardcopy readers for my long web addresses, I am going to spoof this link situation up a notch and simply say: Google for “SAS 9.4 Troubleshooting the SAS Server Tier” and you will find helpful troubleshooting tips for SAS servers.

In case you run into any problems/errors/warnings with the SAS object spawner, here are the log locations:

### **SAS Object Spawner log file locations:**

SAS object spawner log files are located here:

Windows:

**configuration-directory\ObjectSpawner\logs**

UNIX:

**configuration-directory/ObjectSpawner/logs**

And on to the next tier ...

### **Middle Tier**

The middle tier, or web tier, enables access with a web browser using web clients. It is pretty straight forward. Examples for web clients are SAS Studio or SAS Environment Manager.

**Tip:** If you don't know the links to access these clients, you can look at the `instructions.html` file, which is located per default in your configuration directory: `\config_dir\Lev1\Documents`

The middle tier includes the following components:

- SAS Web Application Server and SAS Web Server (http server)
- Cache Locator

- JMS Broker
- SAS Environment Manager

---

### **SAS Web Application Server and SAS Web Server (http server)**

SAS 9.4 has its own web server and web application server. This is good news because it doesn't require the installation of any third-party web app products, such as IBM WebSphere or JBoss Application Server. During the SAS install, the middle tier is installed and configured for you.

The SAS Web Application Server is based on a commercially available third-party software product. The SAS Web Server, an HTTP server, is based on a Pivotal Web Server.

If you look at running processes, you will see *Pivotal Web Server*, which is the SAS Web Server, and *SpringSource tc Run time*, which is the SAS Web Application Server.

Looking at the configuration directory of the SAS Web Application Servers, you will notice folder names such as SASServer1\_1, SASServer2\_1, and so forth. The names SASServer*n* are assigned to certain products in your SAS environment. The configuration directory is: `\sas_config_dir\Lev1\Web\WebAppServer`

Figure 2.6 shows the server assignments. Depending on the products you have licensed, you won't have all the servers and with that, probably no sequential numbering.

**Figure 2.6: Server Assignments**

SASServer 1	Business Intelligence
SASServer 2	Tools
SASServer 3	Performance Management Solutions (General)
SASServer 4	Performance Management Solutions (OLAP)
SASServer 5	Performance Management Solutions (Financial)
SASServer 6	Customer Intelligence Solutions
SASServer 7	Decision Management Solutions
SASServer 8	Risk Management, Fraud & Security Intelligence Solutions
SASServer 9	Performance Management Solutions (Cost & Profitability)
SASServer 10	Performance Management (IT Resource Management)
SASServer 11	Various products (Mid-Tier Support)
SASServer 12	Analytics (Including Visual Analytics & Visual Statistics)
SASServer 13	Data Management Solutions
SASServer 14	Platform Web Services for SAS (Job Scheduler, LSF)
SASServer 15	Data Management Solutions (SAS Data Director)

For details, see SAS Web Application Server Assignments at:

[https://go.documentation.sas.com/?docsetId=biwaag&docsetTarget=n1fojaysjal45on1wio1kp\\_d3u8as.htm&docsetVersion=9.4&locale=en](https://go.documentation.sas.com/?docsetId=biwaag&docsetTarget=n1fojaysjal45on1wio1kp_d3u8as.htm&docsetVersion=9.4&locale=en)

---

## Cache Locator

The SAS Web Application Server uses the cache locator to locate SAS applications that use a data cache for sharing data.

---

## JMS Broker

The JMS Broker is used to send messages between web clients.

---

## SAS Environment Manager

New in SAS 9.4, SAS Environment is a web client that is based on VMware Hyperic, that can be used to monitor and report on your environment. We will talk about this in more detail in Chapter 3.

---

## Java Runtime Environment

There is no need for a JRE in SAS 9.4 as the Java environment is now part of SAS 9.4, compared to prior SAS 9 versions.

A quick note about the Java Runtime Environment, as customers ask about it once in a while. In SAS 9.4, there is no need for a JRE because the Java environment is now part of SAS 9.4, compared to prior SAS 9 versions.

Other web services that you might notice are presented in Table 2.2:

**Table 2.2: SAS Web Services**

SAS Server Name	VMware Hyperic Service
SAS Environment Manager	Apache Tomcat
SAS Environment Manager Agent	HQ Agent
SAS Web Server	Pivotal Web Server
SAS Web Application Server	SpringSource tc Run time
SAS Web Infrastructure Platform Data Server	PostgreSQL

Knowing the services names helps when it comes to monitoring your environment, so you can recognize what process is actually running.

Wait for it – a link reference ... There is a great article called “SAS 9.4 middle tier architecture: need a map?” which I would like to recommend.

<https://blogs.sas.com/content/sgf/2013/09/19/sas-9-4-middle-tier-architecture-need-a-map/>

## **SAS Client Tier**

There is not much to say about this tier: it is where your SAS desktop clients run and where your users access SAS web clients via browser.

---

## **Data Tier**

The data tier “holds” all your data sources. This can be SAS data sets, DBMS tables, multi-dimensional data etc. – simply everything data. SAS offers many options for connecting and accessing data sources, some of which we will cover in Chapter 5, Metadata Library Administration.

---

## **Summary**

We covered a lot of ground in this chapter. The idea is to provide you with basics, to help you understand the SAS architecture. If you know how it works, and what it includes, SAS administration will be much easier. Here are some good resources for general SAS architecture:

Concluding this chapter, I would like to point out one last find, called “Grand Designs: Why It Pays to Think About Technical Architecture Design Before You Act”:

<http://support.sas.com/resources/papers/proceedings13/475-2013.pdf>

When it comes to architecting SAS, or, adding or modifying an already existing SAS environment, be considered and think it through carefully, make sure you have a good plan in place.