

Understanding the Anatomy of a SAS® Deployment: What's in My Server Soup?

Mark Schneider, Donna Bennett, and Connie Robison, SAS Institute Inc., Cary, NC

ABSTRACT

Did you want to better understand the pieces of a SAS metadata-based deployment and where they go and what they do? This paper highlights the major components of a metadata-based SAS deployment of solutions and business intelligence. Using the concepts of topology and tiers as a framework, this paper illustrates the location of key components and how they can be distributed across machines.

INTRODUCTION

Administrators have a great deal of flexibility when defining the topology for a SAS deployment. The SAS® Business Analytics Framework includes many tools to help solve specific business problems. Customers who implement a metadata-based deployment (the SAS Intelligence Platform) will have SAS software, a Java application server, and third-party software. This software framework supports shared business processes, including data management, analytics, and reporting.

This paper describes the components found in typical metadata-based SAS deployments and introduces some common practices for distributing those components across computer resources. End users access user interface components such as SAS® Enterprise Guide® and SAS® Web Report Studio to derive business intelligence from data. Administrators, on the other hand, care more about the underlying infrastructure and content, how those components are distributed across computer systems, and how they interact with available computing resources. They use tools such as SAS® Management Console and the SAS® Deployment Manager to manage the SAS software components.

This paper provides a quick overview of the components (that is, the ingredients) of a SAS deployment and how they blend together. Primary audiences include the following:

- teams who are just starting to think about a SAS deployment
- administrators who will be managing the daily operations of the SAS deployment
- administrators responsible for networking, security, or Web environments who need to understand a bit more about how their areas interact with the SAS deployment
- teams who assume responsibility for a SAS deployment that was created by someone else

READING THE RECIPE

First, we outline some typical topologies, starting with a basic recipe for a single machine and advancing to variations of more complex, distributed configurations. We present three general categories of topologies: simple, basic distributed, and advanced distributed. You can think of these as “starter recipes” that you will spice up to meet specific requirements.

Then we explore the basic ingredients of SAS topologies in more detail: machines, tiers, servers, and third-party software. As administrators, we interact with the set of machines on which we install and configure SAS software. We recognize that the term “server” has many meanings and use the term “server” to refer to software processes. The term “machine” is used to refer to physical or virtual hardware and operating systems on which the software servers run.

Next, we discuss ways you can “spice up” your deployment by choosing options and interactions based on your requirements. We identify key options that administrators can choose, as well as considerations to keep in mind when tailoring a topology for specific needs.

An appendix summarizes key terms and concepts and identifies recommended resources for background information and advanced configuration topics.

The concepts in this paper are covered in much more detail in the administration documents that are part of the SAS Intelligence Platform documentation series.

Designing and implementing advanced deployments is beyond the scope of this paper. Rather, we provide foundational information to help you construct your own custom “recipe” for a SAS deployment. The Recommended Reading section includes references to materials that discuss advanced deployment topics.

RECIPES: COMMON TOPOLOGIES

Topologies describe the basic recipes and proportions used for our “server soup.” A typical deployment of the SAS Intelligence Platform is described as an n-tier architecture, such that SAS servers can be deployed to one or more physical (or logical) machines and accessed via client applications running on desktop computers or via browsers.

Topologies range from simple, single-server machine recipes, which serve as excellent learning examples for new administrators, to complex multiple-machine recipes which distribute computational load and Web traffic for increased availability and throughput. Table 1 highlights some deployment examples within these increasing levels of complexity.

Simple	Basic Distributed	Advanced Distributed
All SAS server components on the same machine	Metadata server and compute server (SAS Application Server) on separate machines	Metadata server on a separate machine and multiple workspace servers on multiple machines; can be optimized for grid and high-performance computing
Might use third-party database provided by SAS	Third-party databases possibly on separate machines	Third-party databases on separate machines
Java application server and SAS servers on the same machine	Java application server on a separate machine	Java application server on a separate machine, possibly optimized with clustering
End-user clients possibly on separate machines	End-user clients on separate machines	End-user clients on separate machines; Java Web Start-enabled clients
Example: SAS® Enterprise Miner™ for Desktop	Example: SAS® Enterprise Miner™	Example: SAS® Marketing Automation using SAS® Enterprise BI Server in a grid environment

Table 1. Examples of Topology Complexity

After we look at a few general topologies, we will drill into the components in more detail and explore the various ways you can customize the mix. The following diagrams are samples of logical views of the three types of topologies and their respective distribution of deployed components. If you have deployed SAS software, you might have seen similar diagrams. Starting with the 9.2 release, SAS delivers diagrams of this type with all standard and custom deployment plans. In these diagrams, each rectangle represents a physical machine, with the client-tier machine serving as a template for multiple end-user workstations.

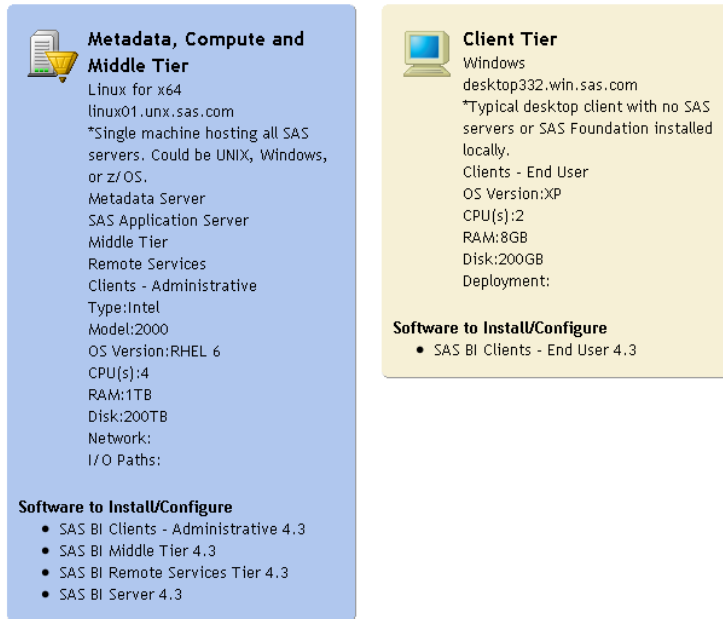


Figure 1. Sample Simple Deployment: Single Server Machine

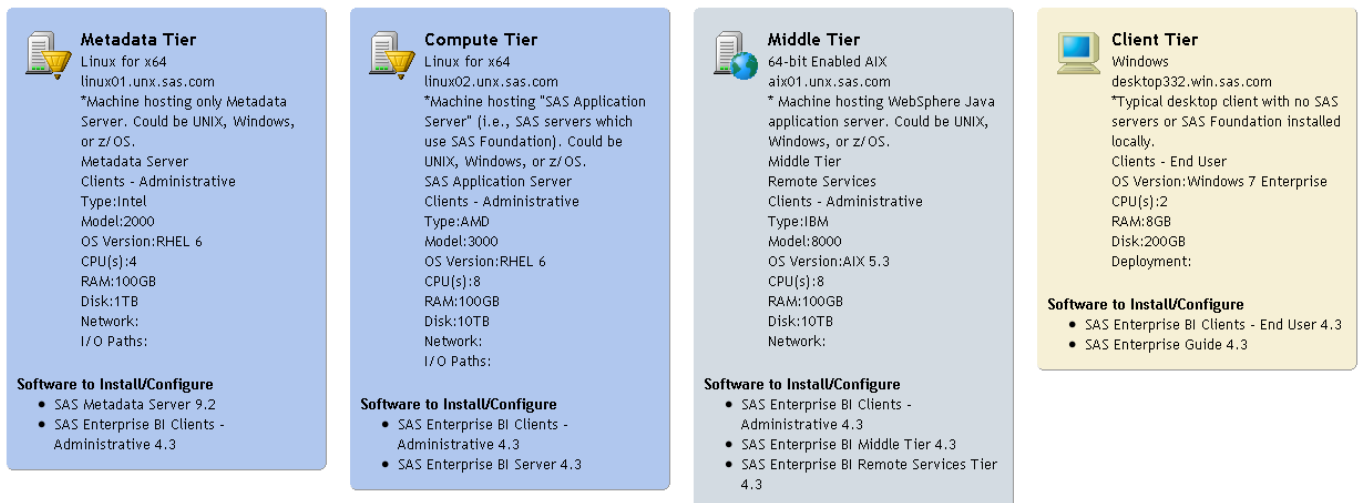


Figure 2. Sample Distributed Deployment: Metadata, Compute, and Middle Tiers on Separate Machines

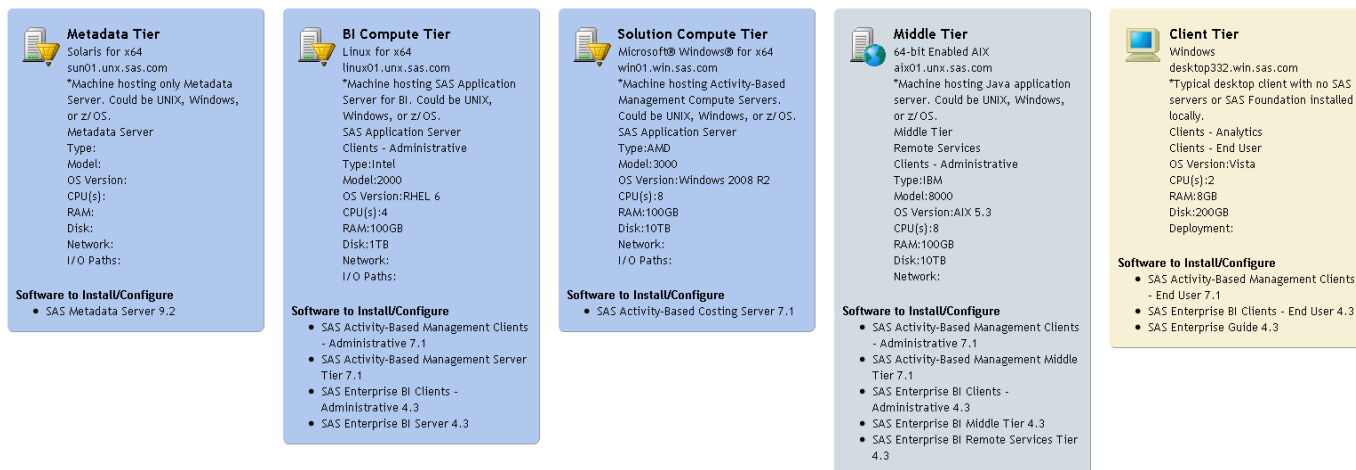


Figure 3. Sample Advanced Distributed Deployment: Dedicated Solution Compute Tier

STOCKING THE PANTRY: TOPOLOGY INGREDIENTS

As shown in the preceding examples, SAS topologies come in many shapes and sizes. The flexibility SAS offers in distributing the software components allows you to customize deployments to your specific user and computing environment. We will now discuss several ingredients you should consider as you concoct a recipe to suit your needs: machines, tiers, servers, and third-party software. In this discussion, we use the term “component” as a generic reference to installable and configurable entities in a SAS topology.

INGREDIENTS: MACHINES

Machines refer to both the physical and virtual hardware on which you will run SAS software. Virtual machines are simply logical partitions of physical machines, providing isolated computing and network resources to one or more users. In all cases, machines are uniquely identified by network host names and run operating systems which support SAS software processes.

Knowing the set of machines hosting SAS software and third-party components is not sufficient to fully describe the topology of a deployment. You should also understand the aggregation of components onto logical tiers and the further breakdown of such tiers into individual SAS servers.

INGREDIENTS: TIERS

We use the term “tier” to describe a logical collection of servers and capabilities that work together to achieve a common purpose. A tier can encompass multiple machines; likewise, a single machine can contain one or more logical tiers. The single server machine topology in Figure 1 shows the metadata, compute, and middle tiers all combined onto a single machine. Conversely, the advanced topology in Figure 3 shows the compute tier split across two physical machines, one supporting BI computing and the other supporting solution-specific computing.

In large deployments, the compute and middle tiers are commonly distributed across multiple machines to provide improved throughput and availability. Administrators often combine the metadata, compute, and middle tiers when setting up development or test environments, and then split them up into corresponding production environments.

While it is possible to have a separate machine for practically every SAS server, such an extreme distribution is not practical to manage and would be highly unusual. On the other hand, one server that is typically assigned to a dedicated machine is the metadata server.

The following tiers are frequently referenced when discussing the topology of a SAS deployment. Examples of components appearing on each tier are shown in Table 2.

- Metadata tier—manages information shared by all components in a deployment; controls authentication and authorization; enables communication among server processes
- Compute tier—also known as the SAS Application Server tier; provides SAS analytic and business processing to requests initiated by the client and middle tiers
- Middle tier—also known as the Web tier; supports Web applications and services
- Data tier—hosts data sources including SAS tables and third-party databases

- Client tier—also known as the end-user or business-user tier; personal workstation hosting user interfaces to the SAS environment, including executable applications (such as SAS Enterprise Guide[®]) and the Web browsers in which SAS Web applications run

Server Contexts

Both the metadata and compute tiers have logical constructs that are referred to as “server contexts.” SAS delivers some default server contexts, and for advanced deployments, site administrators can create additional server contexts. The default server context on the compute tier is typically named “SASApp” and is a logical grouping of the servers listed in the Compute Tier column of Table 2.

The metadata tier contains its own special server context called “SASMeta,” which groups the metadata server itself, along with a workspace server and a DATA step batch server that support administrative functions such as configuration steps run by the SAS[®] Deployment Wizard and metadata repository backups run from SAS Management Console. Access to the SASMeta server context is typically restricted, because these administrative functions require full metadata access as well as the ability to execute operating system calls via the “X” command. Note: Customers who have used the SAS migration tools to move from 9.1.3 to 9.2 might have a “SASMain” server context for the metadata tier.

Metadata Tier	Compute Tier	Middle Tier	Data Tier	Client Tier
SAS [®] Metadata Server	SAS Workspace Server	Java application server (for example, IBM WebSphere Application Server)	SAS data tables	SAS Enterprise Guide
Metadata repositories	SAS Pooled Workspace Server	SAS Web Report Studio (Web application)	Third-party databases	SAS Enterprise Miner client
SAS Management Console	SAS [®] Stored Process Server	SAS [®] Information Delivery Portal (Web application)		SAS [®] Add-In for Microsoft Office
	SAS Object Spawner	Web Infrastructure Platform (includes SAS [®] Content Server)		Web browsers such as Internet Explorer or Firefox
	SAS [®] OLAP Server	Remote Services		SAS [®] Data Integration Studio

Table 2. Examples of Components on Each Tier

INGREDIENTS: SERVERS

The term “server” is used to mean a distinct software process that provides either resources or services to requesting clients, possibly over a network. Many servers can run on a single machine, and servers can collectively support the functional requirements of a tier. Servers come in a variety of flavors, including the following:

- SAS servers (built on the SAS Integrated Object Model (IOM) framework) – execute analytical, reporting, and business processes on behalf of distributed clients. Examples include the following:
 - SAS Metadata Server—supports the previously described metadata tier functionality.
 - SAS Workspace Server—an on-demand SAS execution environment that is instantiated upon client request. Administrators typically assign longer-running stored process programs to this type of server.
 - SAS Pooled Workspace Server—a persistent pool of workspace servers available to clients to query relational information maps.
 - SAS Stored Process Server—a persistent SAS execution environment for stored process programs. Administrators typically assign shorter-running stored process programs to this type of server since multiple users share the same server process.
 - SAS OLAP Server—provides access to multi-dimensional cubes.
- Data Servers—optimize access to a variety of data sources, including the following:

- SAS data sources such as SAS[®] Scalable Performance Data Server[®], OLAP, and SAS/SHARE[®] servers
- External data sources such as Oracle, SQL, and ESRI. (Note: Reading and writing to external data servers might require a companion SAS/ACCESS[®] product on a compute tier machine.)
- Distributed computing SAS servers—distribute workloads across multiple machines and to the appropriate SAS servers. Examples of SAS servers that distribute workloads are the SAS Object Spawner and SAS/CONNECT[®] servers. On grid topologies, SAS[®] Grid Manager helps distribute and monitor the workload.
- Other specialized servers serve unique functional purposes. For example:
 - SAS Content Server—provides access to a WebDAV repository for use by SAS Web applications to store content such as documents, reports, and images. This server is part of the Web Infrastructure Platform—The repository content can be stored directly in the file system or assigned to a database.
 - SAS[®] Table Server—provides scalable, SQL-based data access to integrate data from multiple data sources; supports auditing capabilities of SAS Web Applications
 - SAS DATA Step Batch Server—enables execution of DATA step code in batch mode
 - Third-party scheduling servers, including servers for OS scheduling and Platform Suite for SAS
 - Java application server (also known as Web application server)—JBoss, Oracle WebLogic, or IBM WebSphere J2EE server which hosts SAS Web applications. The Third-Party Software page on support.sas.com provides detailed usage and version information.
 - Web server—provides Web front-end for static HTTP content.

INGREDIENTS: THIRD-PARTY SOFTWARE AND COMMUNICATION PROTOCOLS

SAS deployments rely on third-party software in addition to SAS software. In most cases, third-party software is reinstalled on user machines and customized during the deployment of SAS software.

The SAS components running on each tier rely on established communication protocols to support basic processing and to share information. Although most of the required protocols are not installed or managed by SAS, they need to be accessible to your deployment to enable communication between the SAS components.

Some components are associated with specific tiers, and others need to be available to all machines supporting the deployment. Details about third-party software, including supported versions, are regularly updated on the support.sas.com Web site. The following software components and communication and interface protocols might be needed for your SAS deployment:

- **All Tiers:** The following Java and communication components are needed on one or more tiers of a SAS deployment:
 - Java Runtime Environment (JRE), Java Remote Method Invocation (RMI), and HTTP.
 - TCP/IP, including multicasting, Common Object Request Broker Architecture (CORBA), and SAS Integrated Object Model (IOM).
 - JUnit—optional, but required if you want to use the SAS Deployment Tester to validate the operation of your deployment; installed on machines with SAS Management Console.
 - X11 (UNIX and z/OS deployments)—SAS deployment tools use X Windows technology to support a robust GUI interface. The Third Party Software Reference information on support.sas.com includes references to freely available X11 software.
- **Compute Tier(s)**
 - Platform LSF—scheduling software to coordinate SAS job execution
 - Platform Process Manager and Grid Management Service—support a grid compute environment
 - SAS/ACCESS, as needed to interface to third party data sources
- **Middle Tier**

- Java application server (with J2EE support)—supports execution of SAS Web applications, plus an optional Web server to serve static content. Some sites create multiple managed servers within the Java application server, and these managed servers can be distributed across multiple machines.
- Java Development Kit—provides Java services to the Java application server.
- TCP/IP, HTTP, RMI, and JDBC. SAS services use a multicast address to support some communication.
- **Data Tier(s)**
 - Oracle, DB2, Netezza, Teradata, SQL-Server, and other third-party database software. In addition to storing corporate data, these databases can be used to host some SAS operational structures.
 - Communication between a third party data source and SAS servers might require a SAS/ACCESS license.
- **Client Tier(s)**
 - Microsoft Internet Explorer and Firefox—browser interface to SAS Web applications. Applications accessed via the browser interface are often referred to as “Web clients.”
 - Microsoft Office—Excel, Word, Outlook, and PowerPoint are extended to have direct access to SAS reports and processing when SAS Add-In for Microsoft Office is installed.
 - Microsoft Component Object Model (COM), .NET, and HTTP—leveraged by SAS desktop client software.

SUMMARY OF A WELL-STOCKED PANTRY

Taking stock of what is in your deployment pantry, you will have software, data, and content. Software can include SAS servers, Web applications, third-party software, and user interfaces (clients). We talk about the logical groupings of these components as tiers and the distribution of the tiers onto machines as topology.

Other administrative documents from SAS focus primarily on the categories of components (tiers). Therefore, we include Figure 4 to show a slightly different way of looking at the contents of your pantry.

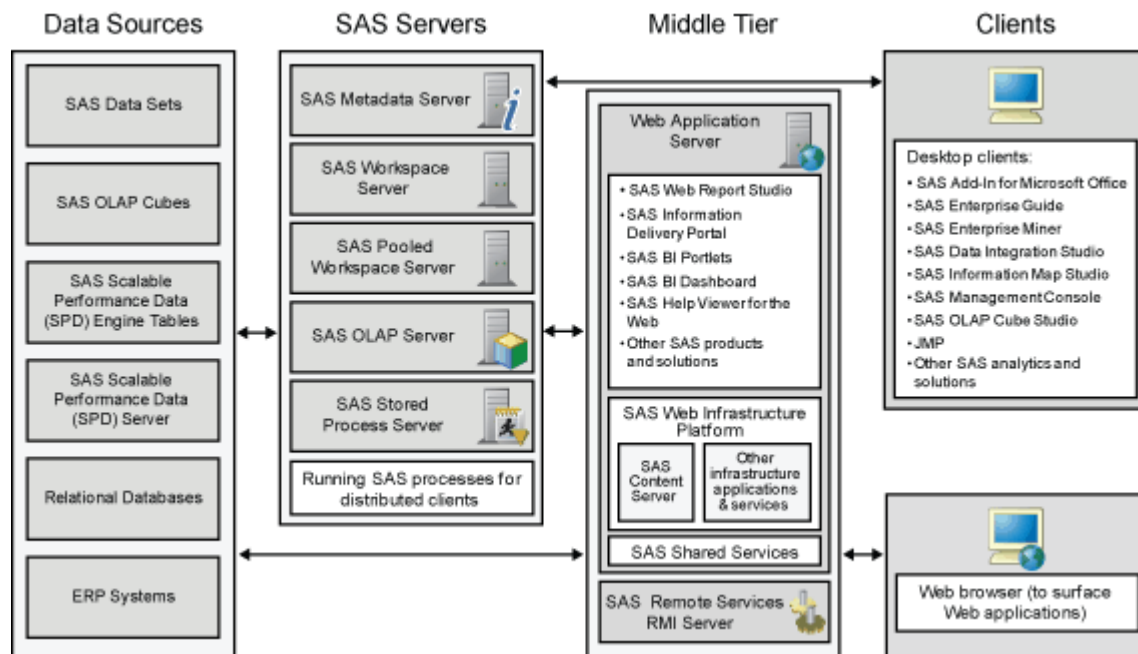


Figure 4. Categories of SAS Components

The logical components of a SAS Business Intelligence deployment are allocated to one or more machines when you define your SAS topology. The deployment plan used as input to the SAS Deployment Wizard determines which SAS software and data components are installed and configured on each machine.

SPICING THINGS UP: DETAILED ADMINISTRATION TOPICS

When you describe the distribution of SAS servers and third-party software across logical tiers and physical machines, you are providing a high-level description of the topology. This description, however, does not identify all of the “moving parts” of a SAS deployment. While users of SAS Enterprise Miner might be satisfied with simply knowing how to connect to a SAS server, users with SAS administration responsibilities need more detailed information. We organize the discussion of such detail into the following four categories:

- SAS Software Depot—central location containing disk copies of installation media
- Installation and configuration files—files that are managed by SAS and are stored in the SASHOME and configuration directories associated with your deployment
- Integration into your computing environment—connection points and corresponding customizations involving network, security, Web, and data administration
- User and administrative content—data, reports, jobs, and other information that users create, modify, or access with SAS software

INSTALLATION AND CONFIGURATION FILES

When you receive software from SAS, either via electronic software download or on physical media, it is delivered in the form of a SAS Software Depot. You deploy software from this depot, and that process creates (or updates) a SASHOME directory and a SAS configuration directory. SAS refers to the directory structure containing installed SAS software files as “SASHOME.” SAS refers to the directory structure containing site-customized and operational files (for example, metadata, custom Web application themes, server scripts, logs, and configuration files) as the configuration directory. Every machine that is part of a SAS deployment has a SASHOME directory tree. Configuration directories exist on most of the machines in a SAS Intelligence Platform deployment, the noted exception being client tier-only machines. Default pathnames for SASHOME directories and configuration directories vary by SAS release and operating system.

SAS Software Depot

When you download software using the SAS® Download Manager, the software is placed in a SAS Software Depot. A single depot maintains disk copies of installation media for all of your orders, optimizing space by storing a single copy of any product which appears in multiple orders. The SAS Deployment Wizard is located at the root of the SAS Software Depot alongside folders that contain license files, third-party support files, various deployment utilities, and the packages from which products are installed and configured.

When you receive an order on physical DVDs from SAS, you run the SAS Deployment Wizard to copy all of the disks to a single SAS Software Depot on your hard drive.

You should store your SAS Software Depot on central network disk space. You can then run the SAS Deployment Wizard on each of your machines directly from this network-accessible depot. You can save time and disk space if you maintain all of your SAS orders in a single depot. You save space by sharing content across orders, and you save download time by downloading only the product content that has not already been downloaded as part of another order.

SASHOME

SASHOME exists on every machine where you install SAS software. Some sites have described this location as “the installation directory.” The installation feature of the SAS Deployment Wizard creates and manages the files in the SASHOME directory. Some very high-level configuration files are located in SASHOME, but administrators should avoid changing these files, opting instead for the server-specific configuration files in the configuration directory. Nonetheless, administrators can make machine-wide changes affecting all users and servers pointing directly and indirectly to SAS software by editing SASHOME configuration files.

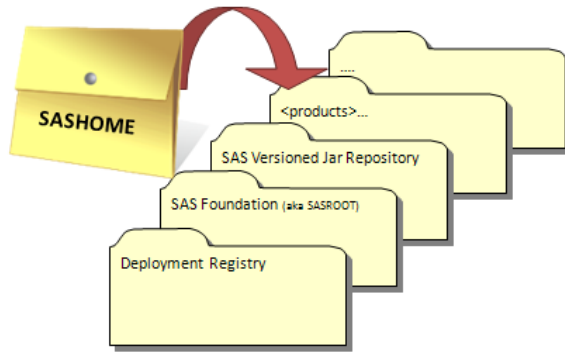


Figure 5: SASHOME Directory

Deployment Registry

Introduced with SAS 9.2, the SAS Deployment Registry is a critical component of every SASHOME. Stored in the “deplymntreg” subdirectory of SASHOME, this data structure keeps track of the SAS software components and versions that are installed on that machine. To view the content of this structure, you can use the View Registry tool, which is documented in the *SAS® Deployment Wizard User’s Guide*. The SAS deployment tools reference the deployment registry when upgrading, applying maintenance, or adding new software.

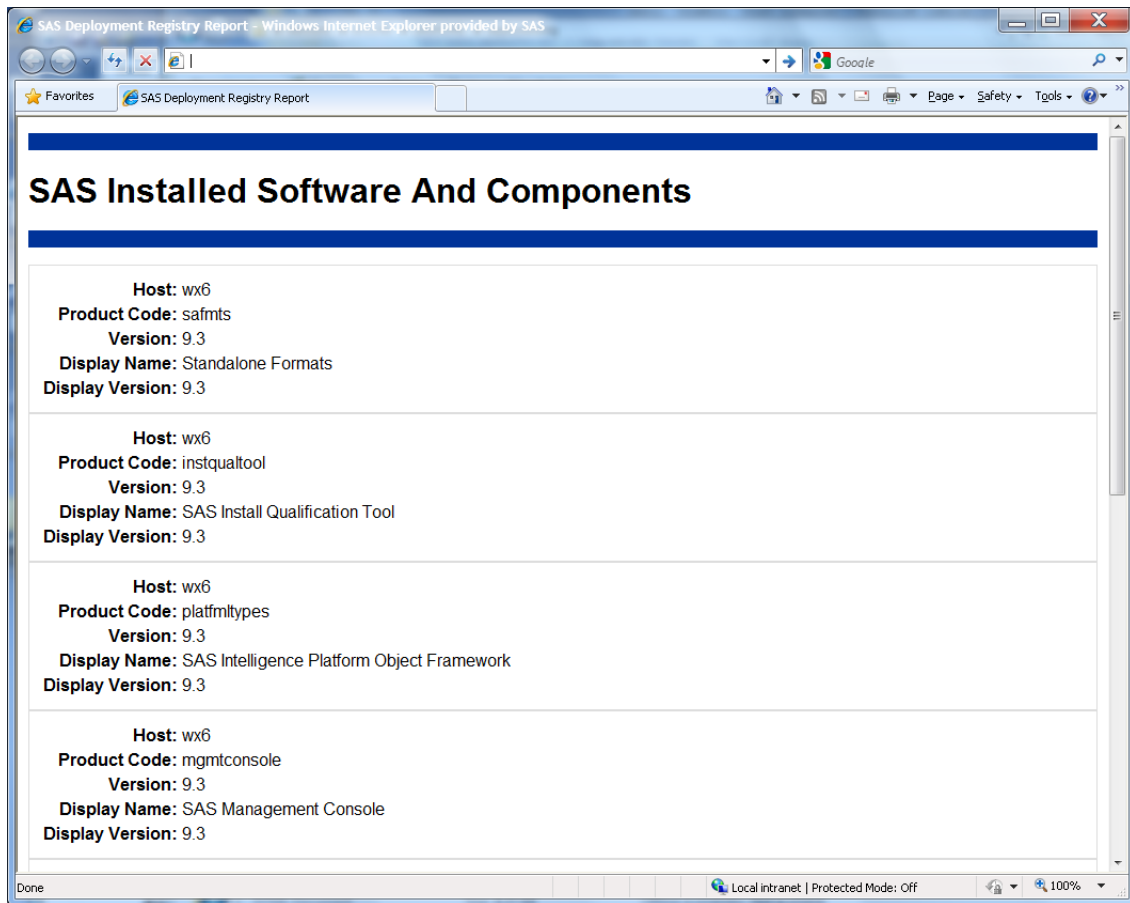


Figure 6: Sample Report Generated by the View Registry Tool Using Information in the Deployment Registry

Configuration Directory

Configuration directories exist on all machines where SAS software is installed, except those which contain only client-tier products. Using SAS deployment tools, you configure your SAS products and define site-specific settings.

The structure and contents of the configuration directory vary depending on the operating system, included tiers, and installed products. Server-specific configuration files are located in server-specific subdirectories of the configuration directory. For example, the MetadataRepositories directory exists only on the machine where the metadata server is installed, and the Web directory exists only on middle-tier machines. You customize your server configuration(s) by modifying “*_usermods” configuration files in these subdirectories. Such customizations can then be preserved when you upgrade to newer versions of SAS software.

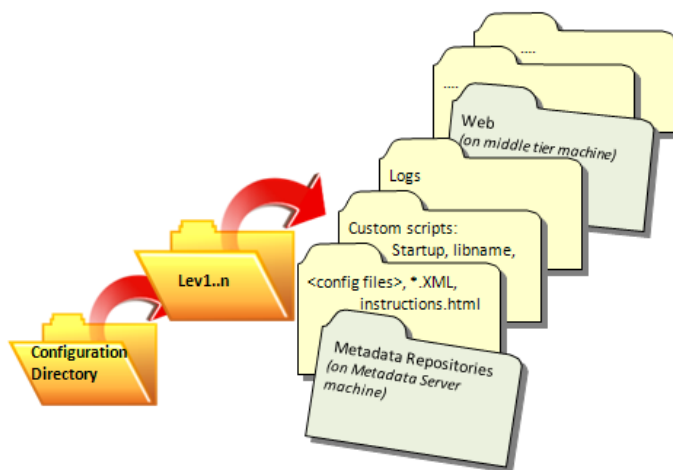


Figure 7: Configuration Directory (Web and MetadataRepositories Directories Do Not Exist on All Machines)

On the metadata tier, the configuration directory contains the metadata repository. On the metadata and compute tiers, the configuration directory also contains the files that support the various servers in your deployment, including start-up scripts, configuration files, autoexec files, and logs. On the middle tier, the configuration directory contains the site-customized copies of SAS Web applications that are ready to be deployed into your Java application server. An option in SAS Deployment Wizard automatically builds and deploys the custom SAS Web applications. Alternatively, you can manually deploy Web applications after running the SAS Deployment Wizard by using your Java application server administration client.

The configuration directory is also the default location for some user content, including stored process code, the SAS Content Server repository, and SAS Table Server relational data.

In addition to registering installed products in the deployment registry (see above), the SAS Deployment Wizard registers configured products in the metadata repository. In this way, the metadata server has access to products configured across different machines and tiers. You run the SAS Deployment Manager to unconfigure any of the products deployed on a given machine, which in turn removes their presence in metadata.

A single SASHOME can point to many different configuration directories (see Figure 8), each being noted in the deployment registry. However a configuration directory can only point to a single SASHOME. SAS supports the concept of configuration “levels” within a configuration directory. Each level (for example, Lev1 or Lev2) represents a separate operational instance of a SAS deployment. Separate configuration directories and/or levels exist for your development, test, and production environments. All of these deployments share the same SASHOME, which means that software changes made to SASHOME affect each of the deployment instances.




Figure 8. Relationship between SASHOME and Configuration Directories

Some sites choose to set up only one configuration level within the SAS configuration directory, and will create a separate SASHOME (and SAS configuration directory) when they need an additional SAS deployment. Both approaches—one or multiple configuration directories associated with a SASHOME—have benefits. The key point is that you need to understand which approach has been used at your site.

INTEGRATING SAS PROCESSES INTO YOUR COMPUTING ENVIRONMENT

SAS software leverages and interacts with a number of components in your computing environment. We have already covered some common third-party software products upon which SAS software relies. In addition, SAS processes depend on other system resources and settings, including data sources, network configurations, security resources, disk space, and memory.

SAS provides a pre-installation checklist that identifies most of the customer-specific information that must be defined when the SAS software is installed and configured. The pre-installation checklist accompanies your deployment plan and deployment diagram. If you used a standard deployment plan, these documents are available on the SAS support site. If you used a custom deployment plan, a customized checklist and diagram should have been part of that delivery bundle.

When you install and configure SAS software, the SAS Deployment Wizard prompts you for custom settings related to data sources, port assignments, users, security options, and disk paths. Some information is created and maintained in the home directory of end users who run SAS software on a desktop machine. Other settings are stored in the metadata repository and various files in the SAS configuration directory.

Data Sources

SAS software products use a variety of data repositories supplied by SAS and by third parties. The SAS Deployment Wizard will prompt you for your desired repository types and locations. If you are using a third-party database management system (DBMS), a corresponding SAS/ACCESS product can be used to bridge communication between SAS processes and the DBMS client software.

Port Assignments and Network Configuration

SAS servers and clients communicate with each other primarily via TCP ports. SAS Web applications running on middle tier machines also use a multicast protocol to broadcast messages over a UDP port. You devote specific TCP and UDP ports to SAS software in order to avoid communication conflicts with other applications. Default port assignments used by SAS are identified in deployment checklists and in administrator documentation.

Users and Security

SAS uses internal (metadata-only) and external (operating system) accounts in its security model. Required system users are noted in your pre-installation checklist.

At a minimum, a SAS deployment must have the following accounts:

- SAS Installer—an external account used to install and configure SAS products
- SAS Spawned Servers (sassrv)—an external account used to run stored process servers and pooled workspace servers
- SAS Administrator (sasadm)—an internal account with unrestricted access to metadata
- SAS Trusted User (sastrust)—an internal account used by SAS servers to communicate securely with one another

The external SAS Installer account is used to deploy all SAS software on a given machine. On UNIX systems, you use this same account to subsequently update or add on to your SAS software. On Windows systems, you can use any account that has administrator rights; it does not have to be the same account.

You create metadata user identities for each individual who requires access to your SAS deployment. You create metadata group identities composed of users who require similar access based on their organization or job function. Metadata-based users and groups do not have to map directly to operating system accounts and groups. In fact, the sasadm and sastrust accounts are generally internal accounts that only exist in metadata. When metadata user identities are linked to external accounts, they can be used to execute workspace server processes, which run under the requesting user's credentials.

The metadata server requires authentication of users to the operating system and uses that information when authorizing access to SAS content. The authentication process can be integrated with existing authentication tools such as LDAP, and can be integrated into Windows single-sign-on capabilities if available at your site.

Metadata-based authorization for accessing content is based on permissions as well as roles. Permissions govern access to metadata objects and associated content (for example, reports, servers, and folders). Roles govern access to specific application features such as menu items, buttons, and plug-ins. In addition, filesystem and other data source security access controls feed into a user's access to information in a SAS environment.

Consult the *SAS® 9.2 Intelligence Platform: Security Administration Guide* for a more thorough description of security.

Disk and Memory

Run-time disk and memory requirements vary based on which SAS products have been deployed. Product-specific system requirements on the support.sas.com Web site describe the minimum disk and memory space you need to initially install and run these products. Actual disk space usage is obviously dependent on user content.

SAS processes frequently generate large blocked I/O disk requests to the operating system, so high I/O throughput rates are important. SAS also uses a large amount of temporary file space when processing data. These temporary files are typically written to the SASWORK directory, which is specified in the server configuration files. Since SAS processing tends to involve large amounts of data on disk, you should invest some time in planning and managing your storage environment. You can find recommendations and FAQs about storage considerations in papers that are listed in the Recommended Reading section.

The metadata server maintains the entire repository in memory, and periodically updates the copy on disk in the configuration directory. The more metadata you store, the more memory you need on the metadata tier machine. It is common to isolate the metadata server on a dedicated machine.

CONTENT: USER AND ADMINISTRATIVE INFORMATION

Many types of content are used within a SAS deployment. We use the term “content” to generally describe a broad range of information that can be created by users or administrators. Content is unique to your site. However, SAS delivers some default administrative content to help with the initial setup of your deployment. In addition, SAS solutions provide some default user content customized for the specific solution domain.

“User content” describes information intended for access by end users. User content can include SAS reports, stored processes, and portlets. Other examples include library references, tables, SAS Data Integration Studio jobs, and SAS Enterprise Miner projects.

“Administrative content” describes information used to define and configure the SAS deployment. Administrative content includes the following:

- definitions of SAS servers within metadata
- specific settings and configuration details needed by SAS applications
- security controls such as access control templates and role and group definitions
- descriptions of SAS resources such as SAS servers and application server contexts
- site customizations, such as themes and branding for Web applications

Most content has two components, both of which are important to consider when you are planning storage layouts or backing up a SAS deployment. The two components of most SAS content are as follows:

- **Description**, which is an object in metadata. For user content, the registration allows end users to access the content without needing to know the specific physical location of the content. SAS security controls are associated with the object in metadata.

- **Information**, which can be contained in one or more data storage repositories such as the SAS Content Server repository, the file system, a third-party database, or the metadata repository. The repository that is used as the data store varies based on the use of the content and the SAS application that manages the content.

SAS Folders enable a logical organization of content, and are based on the description of the content that is registered in metadata. Although not all content appears within the SAS Folders structure, end users and administrators can use SAS Folders to organize content in ways that meet business requirements. In addition, content that appears within a SAS Folder can be exported from and imported into a SAS deployment.

SAS defines default locations for most content, but actual repository locations are very site-specific, and most are defined in response to SAS Deployment Wizard prompts. For example, SAS reports are stored in the SAS Content Server repository. By default, that repository is created within the configuration directory on the middle tier. Administrators can choose to move the content server repository to an alternate location (or machine) and can even choose to put the repository information in a third-party data store. The metadata repository must be co-located with the metadata server. Beyond this restriction, there are few constraints on the specific locations of information repositories.

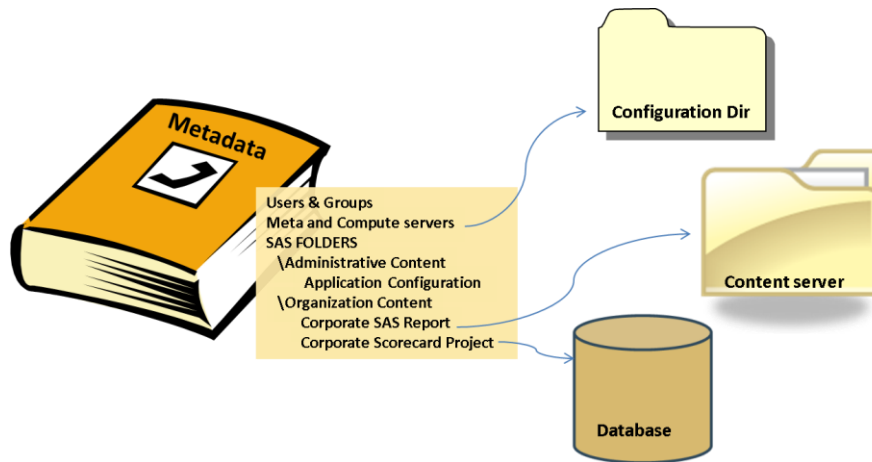


Figure 9. Relationship between Metadata and Information Stores

The SASUSER and SASWORK directories also contain user and administrative content, as well as transient content used during SAS processing. Administrators need to ensure that sufficient space is available for these directories, which are used as follows:

- The SASUSER library contains user-specific customizations that are needed by SAS server processes. By default, a special predefined SASUSER library is located in the home directory of the account that runs the SAS server process. The SASUSER library is often used by SAS DATA step programmers as a convenient place to save permanent data files that are not intended to be published to public libraries.

Note that the stored process and pooled workspace servers are special cases. They reference server-specific directories under the configuration directory rather than the SASUSER library that is associated with a specific end user.
- The SASWORK directory contains temporary files written by SAS servers. SASWORK is defined in the configuration files associated with your SAS servers. Since temporary content written to SASWORK can be large, the directory should have enough free space to avoid jobs failing on the servers. Administrators are encouraged to assign SASWORK to a location that will not compete with other active processes.

PUTTING THE PIECES TOGETHER

The administrator who manages a SAS deployment makes decisions that have an impact on the topology and the details of that deployment. Many of these decisions are made during the installation and configuration tasks of the initial deployment process. Most of the decisions can be revised later.

By understanding the relationships among machines, tiers, and servers you can describe and manage the topology of your SAS deployment. (Note: Although details about administrative tasks are beyond the focus of this paper, it is critical to understand the basic anatomy of a SAS deployment when you are tackling these tasks. The Recommended

Reading appendix identifies some useful resources for more information.) The following are examples of the many situations in which you might need to monitor or change the topology of your current deployment:

- Describing the SAS requirements for security, networking, middle tier, and data to other administrators to prepare for a new or updated deployment, or to manage ongoing changes to an operational deployment.
- Building redundancy into your system to increase availability.
- Scaling for performance.
- Interfacing with security authentication processes such as LDAP.
- Expanding your SAS deployment by doing the following:
 - Adding products. To add products, you will need a new order that includes the new software as well as the latest version of products that are already deployed. You might need to apply updates to the existing products before adding the new software. You also need a SAS deployment plan that incorporates your existing topology plus the products you are adding.
 - Adding servers. If you want to create a new SAS server (for example, a workspace server or stored process server) on a machine which already has a server of the same type, you can simply configure the new server via the Configuration option of the SAS Deployment Wizard. If the machine that is to host the new server does not currently host a server of the same type, you must first run the Installation option of the SAS Deployment Wizard to populate the machine's SASHOME with the necessary server binaries and support files. Then you can run the Configuration option to enable the server in the target environment. In any case, you will need to use a deployment plan and a license file that reflects the server being added.
- Configuring backups. When you are planning your backup strategy, you need to understand where content associated with your SAS deployment is located, when it changes, and how it is related to other content in the deployment. In particular, you need to coordinate backing up content across all machines—especially between the content server repository and the metadata repository.
- Moving your SAS deployment to an alternate set of machines. If the new deployment will have the same topology, you may be able to use SAS migration tools to facilitate the move.

CONCLUSION

A SAS deployment has many moving parts. We have presented the concepts of topology and machines, tiers, and servers as a framework for describing your SAS deployment. You need to understand what those components are and how they fit together as you proceed through each step in the lifecycle of managing a deployment. For example:

- When you are initially planning a deployment, describing the topology can help you select the right hardware and software resources. It also prepares you to work with administrators who manage other aspects of the target environment.
- When you are upgrading or expanding your deployment, understanding the topology helps you ensure that you coordinate tasks across all machines and tiers.
- When you are learning about a SAS deployment that was installed and configured by someone else, describing the topology can help you assess the current state of the software.
- When you are managing the deployment on a day-to-day basis, understanding the topology can help identify key components to monitor and back up.

SAS software provides valuable insight to business users. The concepts and framework presented in this paper aim to provide technical insight into the SAS deployment infrastructure and to provide a starting point for tackling many administrative tasks – from daily management, to troubleshooting, to considering more complex deployments.

TERMINOLOGY

The following list identifies terms that are relevant when describing your SAS deployment. Most of the terms are used within this document.

clients / client software

Clients are software applications that request either resources or services from a server, possibly over a network.

End users can access a SAS deployment using desktop clients and Web clients. Web clients, such as SAS Web Report Studio, run via a browser. Desktop clients, such as SAS Enterprise Miner, are installed on a desktop machine. Desktop clients can be implemented as stand-alone software or can be configured to interact with a shared SAS deployment.

configuration

During the configuration phase of the deployment process, you run tools that create and define the connections and customizations for your SAS deployment.

Information managed by SAS is stored in the metadata repository and in the configuration directory. While many configuration tasks are performed using the SAS Deployment Wizard and the SAS Deployment Manager, some tasks require manual steps.

deployment, SAS deployment

A SAS deployment is an instance of operational SAS software and related components, including SAS components, required third-party software, integration with the host environment, and content.

The deployment process includes installation and configuration tasks and phases. The SAS Deployment Wizard and the SAS Deployment Manager are the primary SAS tools used to support the deployment process.

A SAS deployment is sometimes referred to as a SAS instance or a SAS environment.

deployment plans, deployment diagrams

A deployment plan is an XML file that provides input to the SAS Deployment Wizard. The file contains information about what software should be installed and configured on each machine in a SAS deployment.

Standard deployment plans are delivered with SAS software and support most simple and basic distributed topologies. Custom deployment plans can be created by SAS staff for specialized requirements.

Deployment diagrams provide a pictorial overview of a deployment plan and highlight what SAS software is (or will be) installed on each machine within a SAS deployment.

installation

During the installation phase of the deployment process, you run tools that put SAS software in the appropriate directory locations. The installation phase requires access to a SAS Software Depot.

machine

The term “machine” refers to the resources, real or virtualized, that host a computing environment for software.

In some conversations, the terms “server,” “computer,” and “system” are used to mean the same thing as “machine.”

repository

A repository is a storage container that generally has a defined organization.

Repositories can serve specific purposes. A SAS deployment includes many types of repositories, including metadata repositories and a content server repository. Database systems can be described as repositories that contain specific types of content.

SAS Application Server (or compute server)

Within the SAS Intelligence Platform, a SAS Application Server is a logical entity that defines a collection of one or more SAS servers with specific characteristics.

Users and content can be associated with a SAS Application Server context. Users who submit code from desktop clients such as SAS Enterprise Guide specify the application server context that will be used to execute the code.

SAS Business Analytics Framework

The SAS Business Analytics Framework encompasses world-class solutions, software, and services, enabling the strategic business decisions that optimize performance across your organization. SAS Business Analytics includes a wide range of business solutions, including products for data integration, analytics, and reporting.

SAS Deployment Wizard, SAS Deployment Manager

The SAS Deployment Wizard and the SAS Deployment Manager are deployment tools used to administer your SAS deployment. You use the SAS Deployment Wizard to install, configure, upgrade, and apply

maintenance to SAS products. You also use it to manage SAS Software Depots which contain your orders. You use the SAS Deployment Manager with an existing deployment to perform administrative tasks such as updating passwords, rebuilding Web applications, and changing host names.

The SAS Deployment Wizard is located in the root of the SAS Software Depot. The SAS Deployment Manager is located in your SASHOME directory.

SAS Intelligence Platform

The SAS Intelligence Platform is the technical core of a metadata-based SAS deployment. Documentation about installing, configuring, and managing metadata-based SAS deployments is included in the SAS Intelligence Platform library.

SAS server (or SAS compute server)

The term “SAS server” refers to SAS software that interacts with client software, services, and other servers.

Examples include the SAS Metadata Server, the SAS Workspace Server, the SAS Pooled Workspace Server, the SAS Stored Process Server, and the SAS OLAP Server.

In a SAS business intelligence system, the tier on which the SAS servers execute is sometimes called a “server tier” or a “compute tier.”

SAS Software Depot

The SAS Software Depot is a logical structure that contains disk copies of the installation files needed when you run the SAS deployment tools. You create or update your SAS Software Depot when you receive a software order from SAS via electronic software download or on physical media. You should store your SAS Software Depot on central network disk space.

A single depot can contain multiple SAS software orders. SAS recommends that you retain your depot after completing your initial deployment, and that you use the same depot for updates and additional products.

server

The term “server” refers to software that provides either resources or services to requesting software. The requesting software might be another software process, a Web application, or a desktop client application.

A SAS deployment contains many types of SAS servers and third-party software servers, communicating over a variety of software protocols.

See also SAS server, SAS compute server.

tier

A tier is a logical grouping of components in a SAS deployment. A tier can encompass multiple machines; likewise, a single machine can contain one or more logical tiers.

A metadata-based SAS deployment includes five logical tiers: data, metadata, compute (server), middle, and client.

topology

A topology is a representation of the components of a SAS deployment, the hosting environment and third party software, as well as the physical and logical connections between the components.

In this paper, we describe three general categories of SAS deployment topologies: simple, basic distributed, and advanced distributed. These descriptions highlight the primary components of a SAS deployment – including SAS software and third party software – and the distribution and relationship of those software components across real and virtual machines.

ACKNOWLEDGMENTS

Colleagues from across SAS have contributed to this paper by asking questions, researching answers, and sharing comments and insight. Thanks to the participation from this broad virtual team, this paper provides a comprehensive view of a SAS deployment.

RECOMMENDED READING

Andrews, Brian, Burchett, Gary, and Bennett, Donna. 2010. “The Lowdown on the Download: What’s Really in My SAS® Software Depot and How Did It Get There?” *Proceedings of the SAS Global Forum 2009 Conference*. Available at <http://support.sas.com/resources/papers/proceedings10/318-2010.pdf>.

Chandran, Saravana R. 2010. "The Art of Managing SAS® 9.2 Middle Tier." *Proceedings of the SAS Global Forum 2010 Conference*. Available at: <http://support.sas.com/resources/papers/proceedings10/309-2010.pdf>. This paper contains a good overview of the configuration directories for the middle tier, as well as some troubleshooting tips for the middle tier.

Crevar, Margaret. 2009. "How to Maintain Happy SAS®9 Users." *Proceedings of the SAS Global Forum 2009 Conference*. Available at <http://support.sas.com/resources/papers/proceedings09/310-2009.pdf>.

Crevar, Margaret and Brown, Tony. 2010. "Frequently Asked Questions Regarding Storage Configurations." *Proceedings of the SAS Global Forum 2010 Conference*. Available at <http://support.sas.com/resources/papers/proceedings10/FAQforStorageConfiguration.pdf>.

Crevar, Margaret A., Brown, Tony, and Ihnen, Leigh A. 2007. "Best Practices for Configuring your IO Subsystem for SAS®9 Applications." *Proceedings of the SAS Global Forum 2007 Conference*. Available at <http://support.sas.com/rnd/papers/sqf07/sqf2007-iosubsystem.pdf>.

Kalich, Tanya and Talley, Bubba. 2011. "Troubleshooting Your SAS BI Deployment." *Proceedings of the SAS Global Forum 2011 Conference*.

SAS Institute Inc. 2010a. "Architecture for Business Analytics: A Conceptual Viewpoint." Available at <http://www.sas.com/reg/wp/corp/17871>. This white paper describes the architecture building blocks for empowering strategic business decisions within your existing technology environment. (Note: You must log in to access this paper.)

SAS Institute Inc. 2010b. *SAS® 9.2 Intelligence Platform: Overview, Second Edition*. Cary, NC: SAS Institute Inc. Available at <http://support.sas.com/documentation/onlinedoc/intellplatform/index.html#intell92>.

SAS Institute Inc. 2011a. "Product Documentation: SAS Intelligence Platform." Available at <http://support.sas.com/documentation/onlinedoc/intellplatform/index.html#intell92>.

This Web site provides access to administration documentation for the SAS Intelligence Platform, with dedicated documents for topics that include security administration, SAS system administration, SAS Application Server administration, data administration, Web application administration, grid computing, logging, and scheduling. The following are links to specific documentation topics related to subject matter covered in this paper:

- SAS configuration directory:
<http://support.sas.com/documentation/cdl/en/bisag/64088/HTML/default/viewer.htm#a003181764.htm>
- Contents of the Lev1 directory:
<http://support.sas.com/documentation/cdl/en/bisag/64088/HTML/default/viewer.htm#a003181764.htm>
- Configuration files for SAS servers:
<http://support.sas.com/documentation/cdl/en/bisag/64088/HTML/default/viewer.htm#a003147155.htm>
- Connecting to common data sources:
<http://support.sas.com/documentation/cdl/en/bidsag/61236/HTML/default/a003112367.htm>
- Default ports and protocols:
<http://support.sas.com/documentation/cdl/en/bisag/64088/HTML/default/viewer.htm#a003141390.htm>
- Administering logging for SAS Servers:
<http://support.sas.com/documentation/cdl/en/bisag/64088/HTML/default/viewer.htm#a003286197.htm>
- Log files for desktop applications:
<http://support.sas.com/documentation/cdl/en/bidaag/61231/HTML/default/viewer.htm#a003314426.htm>

SAS Institute Inc. 2011b. "SAS Training: Learning Paths by Job Role." Available at <http://support.sas.com/training/us/paths/index.html>. This Web site includes information about SAS administration training offerings.

SAS Institute Inc. 2011c. "Standard Deployment Plans." Available at http://www.sas.com/apps/demos/downloads/samplePlans_PROD_sysdep.jsp?packageID=000501. This Web site provides standard deployment plans and related checklists and diagrams.

SAS Institute Inc. 2011d. "Third Party Software Requirements for Use with SAS® Products." Available at

<http://support.sas.com/resources/thirdpartysupport/> . This Web site provides reference notes for configuring and customizing middle-tier components, lists of supported versions of third-party software components, links to external sites, and tips for managing third-party software.

CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the authors at:

Mark Schneider
SAS Institute Inc.
100 SAS Campus Drive
Cary, NC 27516
919-677-8000
Mark.Schneider@sas.com

Donna Bennett
SAS Institute Inc.
100 SAS Campus Drive
Cary, NC 27516
919-677-8000
Donna.Bennett@sas.com

Connie Robison
SAS Institute Inc.
100 SAS Campus Drive
Cary, NC 27516
919-677-8000
Connie.Robison@sas.com

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.