

Paper 225-29

Top Deployment Considerations for the SAS® 9.1 Intelligence Architecture

Pat Bostic, SAS Institute, Cary, NC
Clarke Thacher, SAS Institute, Cary, NC
Dave Crow, SAS Institute, Cary, NC
John Roth, SAS Institute, Cary, NC

ABSTRACT

The SAS 9.1 Intelligence Architecture is a multi-tiered platform consisting of components residing on client desktops, middle-tier web servers, and enterprise-class back-end servers. These components can be spread across multiple systems running on a variety of hardware and integrated with a wide range of third-party products. This multi-tiered design offers tremendous flexibility and configurability but requires careful planning and deployment. In this presentation, platform managers from the Host Systems Division present the top considerations for preparing to deploy the SAS 9.1 Intelligence Architecture throughout your organization.

INTRODUCTION

SAS 9.1 introduces big changes in the way a SAS release is deployed throughout an enterprise. Early in the development process, several groups in SAS R&D realized the importance of installing, configuring and deploying the new infrastructure properly. This resulted in complementary initiatives from several groups.

These projects have different names: Customer Experience Testing (CET), Total User Experience (TUE) and Host User Experience (HUE), but they have a common goal: To improve the customer experience in deploying products based on the SAS 9.1 Intelligence Architecture. In this paper, we will describe some of the lessons that we in Host Systems R&D have learned from our work on HUE, where we focused on deploying the SAS 9.1 Intelligence Architecture in a heterogeneous environment with the clients, mid-tier, and server systems residing on different platforms.

THIS IS NOT YOUR FATHER'S SAS

"Back in the old days" putting up a new release of SAS was a pretty simple matter. You put the CD (or tape or floppy disks) into the drive on your machine and installed. In less than an hour, you were up and running. Your only options for client/server were SAS/Share® and SAS/Connect®. The Internet was a dream of folks in some research labs and universities (how many people remember BITNET?). People accessed their systems through terminals like the IBM 3270 and VT100 or card readers. SAS was there, helping you with powerful reporting and analytics.

Things have changed a lot since then. We are expected to have up-to-the minute intelligence about our enterprise distributed throughout the enterprise. The volumes of data have grown enormously and the old ways of working are not sufficient for today's competitive environment. The SAS 9.1 Intelligence Architecture was created to meet this challenge.

The SAS 9.1 Intelligence Architecture provides a suite of servers and services that deliver *The Power to Know*® throughout the enterprise. Applications developed using the SAS 9.1 Intelligence Architecture can be deployed in multi-tier environments where web browsers, application servers and SAS® Foundation Servers operate to create a dynamic and powerful environment to meet today's requirements.

SAS FOUNDATION SERVERS

The servers in the SAS 9.1 Intelligence Architecture include the following:

- SAS® Metadata Server The SAS Metadata Server enables centralized, enterprise-wide metadata delivery and management: One metadata server provides metadata to SAS applications across the enterprise.
- SAS® OLAP Server The SAS OLAP Server delivers pre-summarized "cubes" of data to OLAP clients such as SAS® Enterprise Guide® using OLE DB for OLAP. The SAS OLAP Server is a multidimensional database server that is designed to reduce the load on traditional back-end storage systems by delivering different summarized views of data to business intelligence applications, irrespective of the amount of data underlying these summaries.
- SAS® Stored Process Server The SAS Stored Process Server executes and delivers results from SAS Stored Processes in a multi-client environment. A SAS Stored Process is a SAS program that can be called through the SAS Stored Process Server. The SAS Stored Process Server enables users at client workstations to execute parameterized SAS programs without having to know the SAS language.
- SAS® Workspace Server The SAS Workspace Server provides clients such as SAS Enterprise Guide or the SAS® Add-in for Microsoft Office access to the full SAS programming environment. Users at client workstations can then use their knowledge of SAS programming environment to develop ad hoc reports.

SAS® FOUNDATION SERVICES

This suite of Java-based APIs provides core middleware infrastructure services including user authentication, profile management, session management, activity logging, metadata and content repository access, and connection management. SAS Foundation Services provide extension services to assist in information publishing, event management, and SAS Stored Process execution.

SAS® APPLICATION SERVICES

SAS Application Services provide business-oriented query and reporting services to calling clients. By using a business metadata layer and a universal report definition, SAS® Query and Reporting Services provide a solid foundation for enterprise reporting and application development. Java and COM-based interfaces to SAS Application Services surface to clients the functionality provided by SAS Query and Reporting Services. SAS Application Services can also be used by application developers to provide custom business intelligence capabilities within their solutions.

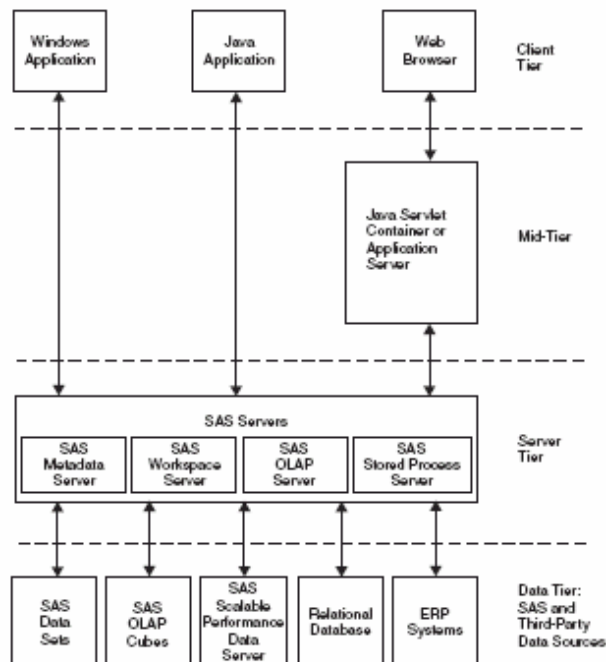
YOUR OLD FRIENDS ARE STILL THERE

You may be thinking “What do I have to do to keep my old stuff working?” Don’t worry: Your old friends are still there. If you only want to migrate existing applications from an earlier SAS release you will find that SAS 9.1 requires minimal effort. For those few changes that are required, excellent tools have been developed to facilitate this migration.

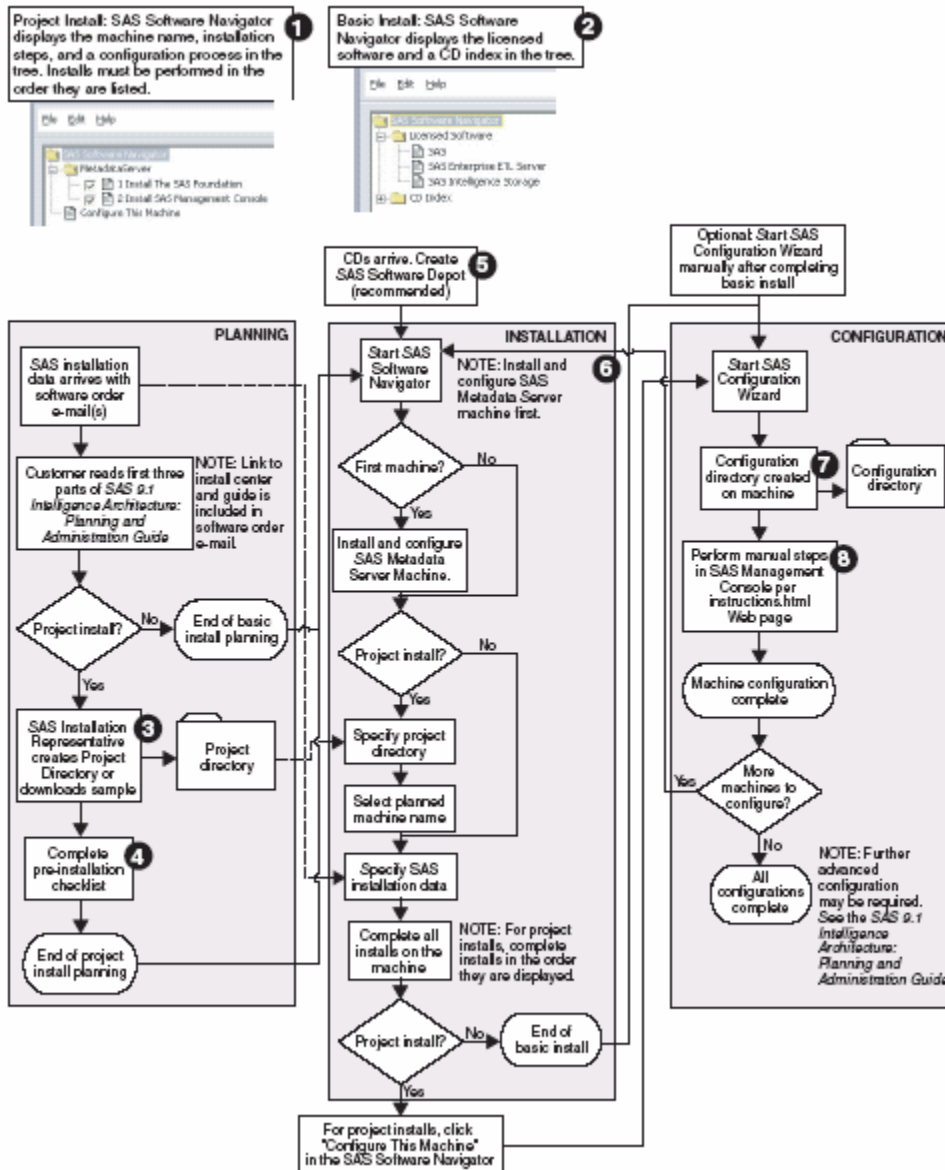
When you do start to utilize the new features in SAS 9.1, you will find that what you already know about using SAS will apply directly to the new environment. The SAS 9.1 Intelligence Architecture is built on the very solid SAS® Foundation that you’ve grown to depend on over the years. The SAS Foundation is the source of core services for data processing, statistical, and analytical power. SAS 9.1 is supported on industry-leading 32-bit and 64-bit operating environments. Parallel processing enhancements have been added to several products including Base SAS® Software (procedures MEANS, REPORT, SORT, SQL, SUMMARY, TABULATE), SAS/STAT Software (procedures GLM, LOESS, REG, ROBUSTREG), SAS/SHARE and SAS® Enterprise Mine™ (procedures DMINE, DMREG). Threaded I/O capabilities have been added for indexing as well as several SAS/ACCESS® engines to DBMS servers (DB2, Oracle, Sybase, etc.). Your investment in the SAS Foundation continues to provide substantial dividends. The SAS 9.1 Intelligence Architecture provides new ways to realize the value of and build upon your investment.

MAKE A PLAN AND FOLLOW IT

The SAS 9.1 Intelligence Architecture is a multi-tier environment.



The deployment of your full SAS 9.1 Intelligence Architecture suite can be overwhelming, involving multiple machines and operating systems. Installation and configuration of all the components can be complicated. To avoid confusion and streamline the process, you should develop a plan with details about your unique system configuration (machines, software and order of installation). The result of careful planning will be a planning file to be used during the installation and configuration process. To help identify the steps you should follow, and the decisions you should make, the planning and administration guide provides a flowchart to illustrate the steps involved in setting up a system:



For detailed information on how to perform the tasks illustrated, see the *SAS 9.1 Intelligence Architecture: Planning and Administration Guide*. The time and effort spent in paying careful attention to detail while planning for system deployment will be rewarded many times over.

OH, WHAT A TANGLED WEB WE WEAVE: PROTOCOLS, APPLICATION SERVERS, AND THE SAS 9.1 INTELLIGENCE ARCHITECTURE

With SAS 9.1, users will find that the features, power, and flexibility offered by SAS can now be readily used over the World Wide Web. To achieve this integration, SAS software itself has been structured to fit within a multi-tier arrangement commonly used by web-based applications, and it uses industry-standard protocols to communicate between tiers. By using industry standards, SAS 9.1 can work smoothly with infrastructure already in place to support the Web, while offering users a familiar, well-understood interface.

In the multi-tier model, the client tier supplies the user interface, while the "back end" runs the SAS Foundation Servers (stored process servers, OLAP servers, etc.) to access data and perform traditional SAS analysis. The mid-tier resides between these two layers providing business logic integration. It receives requests for information from the client tier, and contacts the back end to satisfy these requests. The mid-tier packages the information received from the back end in a suitable format and returns it to the client tier for presentation to the user. This presentation is

not restricted to traditional SAS reports and graphs, which is part of the power of using JAVA in the client workstation.

In SAS 9.1, the mid-tier utilizes an application server and communicates with the client tier using two protocols, Hypertext Transfer Protocol (HTTP), and Web-based Distributed Authoring and Versioning (WebDAV). HTTP is the common language used by all web browsers and servers. WebDAV is an extension of HTTP which adds the ability to collaboratively manage and edit files residing on a web server.

Application servers provide Java 2 Enterprise Edition (J2EE) standard services. The J2EE standard specifies a long list of features, but the most frequently used are Java Servlets, Java Server Pages (JSPs), and Enterprise Java Beans (EJBs). A number of vendors supply application servers. SAS recommends either the Apache/Tomcat combination or WebLogic from BEA. Apache/Tomcat is the reference standard for Servlet and JSP technology. It is available at no cost from the Apache Group. Although Apache/Tomcat works well for development projects and smaller deployments, WebLogic is a better choice for enterprise-class applications. Furthermore, SAS 9.1 Solutions that are built to run on SAS 9.1 require EJBs and other features supported in WebLogic but not supported in Apache/Tomcat. Another very popular application server, IBM's WebSphere, also provides the required features and will soon be supported.

THE SAS METADATA SERVER IS THE CENTER OF YOUR BI WORLD

As we have shown, the SAS 9.1 Intelligence Architecture is composed of many servers, services and clients. The SAS Metadata Server is the glue that binds these elements into a single powerful solution. Integrated metadata (information about data sources, how it was derived, business rules and access authorizations) is crucial for producing accurate, consistent information. By storing metadata from all applications in an open, centralized and integrated repository, data changes only need to be documented in one place, there are fewer systems to support, and business users can count on high-quality, consistent information. A single version of the truth is available to all, and better use of staff time lowers the total cost of ownership for IT infrastructures.

The SAS Metadata Server is a multi-user software server that surfaces metadata from one or more repositories to applications via the SAS® Open Metadata Architecture. With the ability to import, export, reconcile and update metadata, and document those actions, the server manages technical, process and administrative metadata across all applications. It adheres to the Common Warehouse Metamodel standard as defined by the Object Management Group and supports interoperability with third-party products.

YOU CAN PUT IT ON A LAPTOP, BUT YOU SHOULDN'T (PERFORMANCE MATTERS)

The SAS 9.1 Intelligence Architecture is "simple" enough to deploy on a single laptop, but that is not the most optimal configuration to run an enterprise. Indeed, you will have specific performance requirements. Consider the number of users that will be exercising the system. Establish workload scenarios and expected response times. Anticipate peak load. Armed with this type of information, choose appropriate hardware, operating system and application server environments to meet these specific requirements. The information gathered will provide direction for the number of processors, amount of memory, and disk space. The choices will influence the robustness and reliability of the overall system.

Because the metadata server plays such a central role in the SAS 9.1 Intelligence Architecture, its performance is critical to the overall performance of these multi-tier applications. The metadata server was designed with an in-memory database to provide quick response time to all of the components of the SAS 9.1 Intelligence Architecture. Our experience has shown that the SAS Metadata Server performs best if it runs on a system that has sufficient disk space, CPU, and memory resources to meet the demands of the SAS 9.1 Intelligence Architecture. Ideally, the SAS Metadata Server should not have to compete for resources with other applications on the system. This can be accomplished by running it on a physically separate system or in a separate logical partition.

Another general rule of thumb is to put the servers close to the data. For example, OLAP cubes are a powerful tool for summarizing data so that ad hoc queries can be processed efficiently. However, in order to build the cube, all of the underlying data must be processed, and you will probably find that you don't want to transfer all that data across your company network.

SECURITY IS NOT JUST A WARM BLANKET

OVERVIEW

The SAS 9.1 Intelligence Architecture provides powerful tools to extract information from large volumes of key business data. Securing access to these tools, and to the information that they extract is an important consideration during planning and implementation. The **SAS® 9.1 Intelligence Architecture: Planning and**

Administration Guide devotes over 60 pages to security issues. That may seem daunting at first, but this book is truly your friend. The sections on security take you from initial considerations of your business requirements, through how those requirements can be met using the tools provided by the SAS products, all the way to some examples of specific implementation. At each step along the way, it includes suggestions for how to incorporate your own enterprise's security goals into the implementation.

The security implementation for the SAS 9.1 Intelligence Architecture acts to complement and extend the security features provided by the underlying host and database security systems. It is not a replacement; SAS uses the host security systems to implement the necessary security features.

SOME KEY CONCEPTS

The planning and administration guide discusses these concepts and many others in detail. Here are some key concepts to understanding the security environment.

Single-signon: The SAS 9.1 Intelligence Architecture is designed around the concept of a single-signon that is managed by the metadata server. A client identifies itself to the metadata server, and any subsequent servers that must be accessed are authenticated using credentials that have previously been stored in the metadata repository. As long as all credentials are available and valid, a user should not see any subsequent prompts for login information.

Use of Host authentication: By default, servers use host authentication. For example, if a SAS Stored Process Server is running on an AIX system, the application that requests service must provide login credentials that are valid on the AIX system. This can either be a userid that is associated with the individual, or a "pooled" userid that is associated with a group of users.

Use of LDAP or Active Directory for Metadata Server authentication: Like all servers, the SAS Metadata Server uses host authentication by default, which means that any user signing on to the metadata server must have a userid on the host where the metadata server runs. Alternatively, you can direct the metadata server to use an LDAP or Microsoft Active Directory to do authentication.

Application Server Authentication: When using a mid-tier application server, it is possible to use the application server authentication in place of the user authentication provided by the SAS application. In this mode of operation, the client does a signon to the application server (e.g. WebLogic). Once this signon is done, there is no requirement for additional authorization by the metadata server.

Updating the SAS Metadata Repository

Normally, changes to the metadata repository to implement the security environment are made using the SAS[®] Management Console. For bulk changes, such as adding a large number of users to the metadata, it is also possible to use SAS to avoid the manual overhead often associated with such changes.

Authorization Layers

The Metadata authorization layer can be used to specify access controls on metadata objects that represent computing resources such as OLAP cubes, variables in a SAS data library and columns within a database table. When accessing these resources via SAS, this provides the ability to impose access restrictions at a finer granularity than is allowed by the operating system or database access controls. This Metadata authorization layer and the rules for access inheritance are powerful tools in designing an efficient, maintainable framework of access controls.

User Groups

User groups should be used to minimize the complexity of the security environment. Where possible, authorization should be granted to groups rather than individuals according to job function so that individuals can change jobs without requiring massive updates to the security authorizations.

Getting Started

We recommend that you first create a prototype system that is similar to the system you intend to deploy in production. If you intend to deploy a heterogeneous environment in which the various tiers are on different platforms, you should also consider making this prototype system heterogeneous, as there are security considerations that are affected by the mix of host platforms that are involved.

In many cases, the prototype deployment can be done with a simple security system such as the one described in the section "A Minimal Security Environment" within the "Implementing Security" chapter of the planning and administration guide. This will enable you to get familiar with the terms, tools, and processes needed to design your production security environment. Spending a little extra time learning how the SAS 9.1 Intelligence Architecture fits

into your existing environment will save time and avoid problems.

Some notes of caution

For this prototype deployment, stay with the configuration produced by the SAS[®] Configuration Wizard as much as possible. Deviations from this configuration may cause unexpected problems that are difficult to diagnose. We believe that most customers will be able to use the wizard's configuration for both the prototype and the production environment. However, if you find it necessary to tailor this configuration, by all means, do it on the prototype system first.

The Next Steps

The detailed plan for the security environment requires knowledge of the specific systems that you will use to host the SAS Foundation Servers, the mid-tier application servers, and the clients. This plan must be developed in parallel with the specifications of the application environment.

Study the security sections of the *SAS 9.1 Intelligence Architecture: Planning and Administration Guide*. Remember, this book is your friend, pay attention to what it says.

It is important to understand the enterprise standards and the legal requirements that your security system must meet. We recommend that you involve the various groups that are responsible for your IT security from the outset. For example, your network security staff may require that your servers be behind firewalls in order to protect them against network attacks from outside the company. The application design may also have requirements that certain application servers be outside those same firewalls. These issues are not explicitly covered in the planning and administration guide, but they may well affect how you implement your solution.

ALL USERS ARE EQUAL, BUT SOME ARE MORE EQUAL THAN OTHERS: THE REQUIRED USER IDS OF THE SAS 9.1 INTELLIGENCE ARCHITECTURE

OVERVIEW

Before you begin to configure the SAS 9.1 Intelligence Architecture in your enterprise, you will need to define up to 7 new user ids. When using Unix or z/OS as server platforms, you will also be asked to define a security group in the host's security system. Each of these user ids will play a key and distinctive role in the operation of the SAS 9.1 Intelligence Architecture on your network. It is important that you do not attempt to merge these roles into common user ids or your system may not work properly. The names used in this document are only examples, and you may choose other names for the group and user ids according to your enterprise policy.

THE SECURITY GROUP

When running SAS Foundation Servers on Unix and z/OS platforms you will be asked to create a security group named **SAS**. This security group will be used to control access to the directory on which the servers run. Sensitive information is contained in this directory structure, so it must be protected accordingly.

THE REQUIRED USER IDS

The **sas** (**sas**) user id will be used to install and configure all SAS software. Make the **SAS** group this user's default. We recommend that you install and configure all SAS software under this user id. By default this user will be the owner of the OLAP and Metadata servers, the SAS[®] Object Spawner, the SAS/Connect Spawner and the SAS/SHARE server. This user will own the configuration directory structure that these servers use. The configuration directory is protected such that only this ID has access to most of the directory hierarchy. Installing under one user ID and running the servers under a different user ID will cause server failures. This account is highly privileged in the SAS 9.1 Intelligence Architecture and should be protected accordingly.

The SAS Administrator (**sasadm**) user id should be created on your metadata server machine. This user has privileges to manage user accounts in metadata and administer the metadata server. The SAS Administrator has unrestricted access to the metadata and this user ID should be protected accordingly. This ID should only be used for the SAS Management Console application.

The SAS General Server (**sassrv**) user id should be created on your server machines. You must make the **SAS** security group its default group. This account will be used by the object spawner to launch stored process servers. This account will need access to any OS resources required to run stored processes. The default configuration of the SAS 9.1 Intelligence Architecture creates this single account for load balanced, stored process server usage. Additional server accounts can be created to give different levels of access as required.

The SAS Guest (**sasguest**) user id will need to be created on your metadata server machine. This user is a generic user account and should have the lowest level of security privileges. This user id is used by the SAS[®] Information Delivery Portal to log users into the public kiosk area.

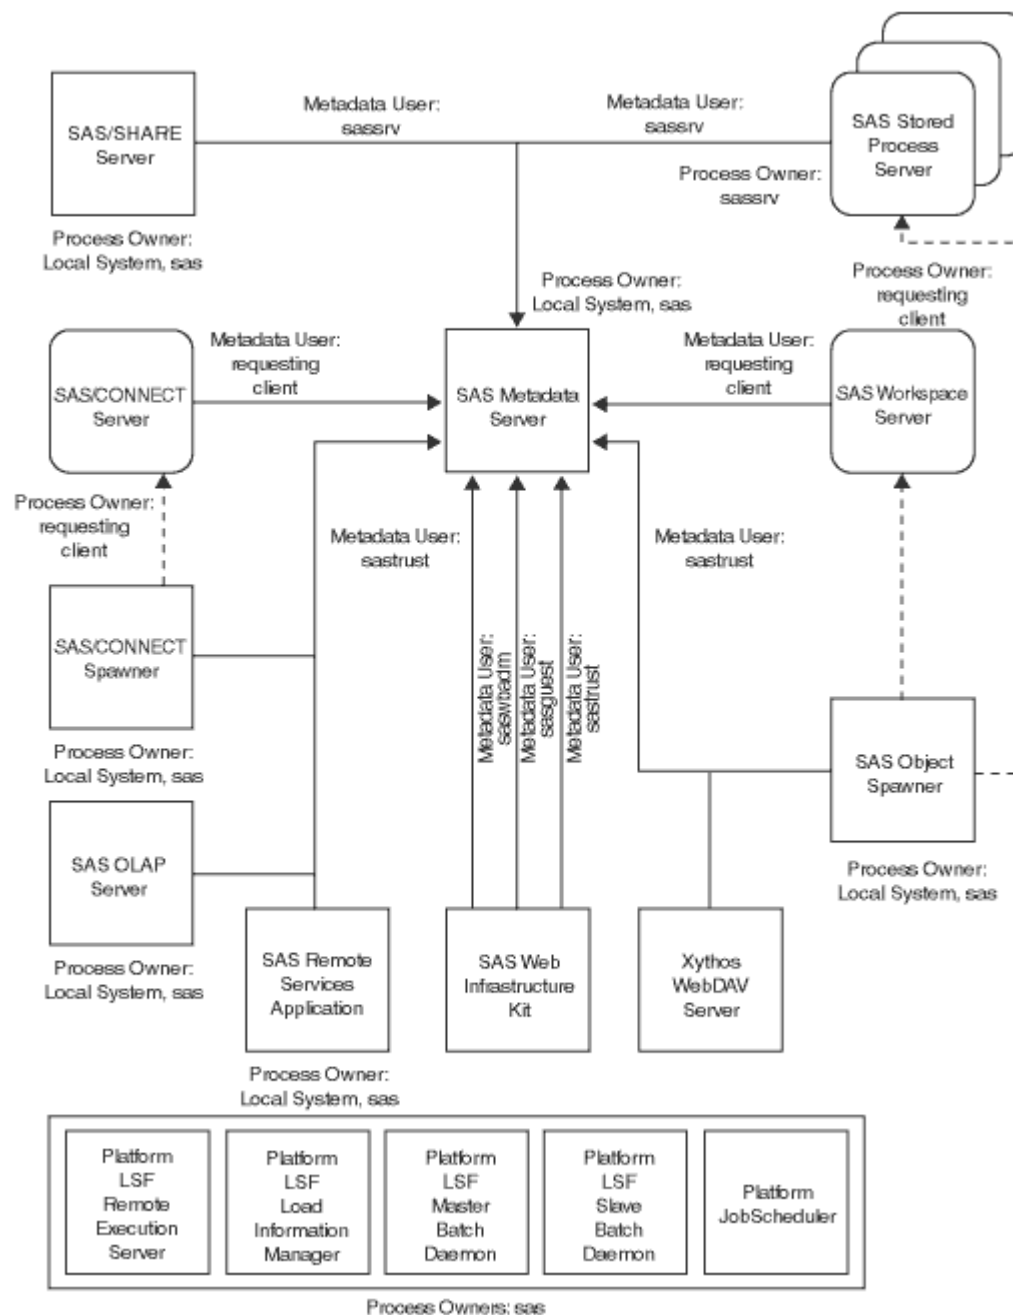
The SAS Trusted User (**sastrust**) will be created on the metadata server machine. Because this user ID is a trusted ID, SAS servers such as the OLAP server and mid-tier applications can use this ID to impersonate authenticated clients on the metadata server; that is, the servers can communicate with the metadata server on behalf of the clients. This is a highly privileged account and should be protected accordingly.

The SAS Demo User (**sasdemo**) will be created on your metadata server machine. This user has permission to demonstrate the SAS software you have installed and to verify the configuration.

The SAS Web Administrator (**saswbadm**) will be created on your metadata server machine. You need only create an account for the Web administrator if you will be installing web applications, such as SAS® Web Report Studio, SAS® Web Report Viewer, or SAS Information Delivery Portal. This user has permission to administer the SAS Web infrastructure.

HOW THE ACCOUNTS ARE USED BY THE SAS SERVERS

The figure below depicts how the accounts you create during pre-installation are used by the SAS servers in an intelligence system. The figure shows who owns each server process and which account each server uses to communicate with the metadata server.



BUILD YOUR BI STRUCTURE ON SOLID DATA

Inaccurate, outdated and incomplete data is the single biggest obstacle to successful implementation of a SAS 9.1 Intelligence Solution. A recent European survey revealed that 66 percent of European organizations believe that data quality and integration issues are impacting company profitability. Multiple platforms, various data formats with external data combined represent the typical corporate data environment. The inability to cleanse data beyond the typical name and address requires considerable effort to understand data elements and interrelationships, vastly increasing the time and cost for deployment. Data usually enters the system from multiple sources, and even though much of this is done with automated processes used by suppliers and customers, data entered is rarely represented identically across the systems.

SAS® ETL Server provides an interactive development environment with a single point of control for managing your processes to Extract, Translate, and Load “cleansed” data. The SAS ETL Server uses the SAS Metadata Server to describe the underlying data warehouse that is being created. This allows you to leverage the work done in defining and executing the ETL process across all of the applications that are built on that data.

CONCLUSION: COME ON IN, THE WATER’S FINE

As we mentioned in the introduction, this paper is an outgrowth of the HUE project within R&D. The people on the project had no previous experience with this new SAS architecture, and we approached it as we expect you the customer will approach it. Overall, we were very impressed with the power of these new tools, and the ease with which we were able to deploy applications that would have been incredibly difficult prior to SAS 9.1. In this paper, we have shown some things that must be considered for a successful deployment of the SAS 9.1 Intelligence Architecture. We believe that these suggestions will help you to deploy an intelligence architecture that can deliver powerful and timely intelligence across your enterprise.

REFERENCES

The SAS® 9.1 Intelligence Architecture: Planning and Administration Guide along with other reference documentation can be found at <http://support.sas.com/documentation/configuration/91admin.html>

The Common Warehouse Metamodel is described at the OMG website: <http://www.omg.org/cwm>.

In addition to these references, you will also want to become familiar with the various communities on the web sites at www.sas.com, as well as the tech support web sites. As with any SAS product, SAS stands ready to provide assistance in any planning, implementation or deployment projects that you have.

CONTACT INFORMATION

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

Pat Bostic, Dave Crow, Clarke Thacher, John Roth, Arthur Hunt, Phillip Smith

SAS Institute

SAS Campus Drive

Cary, NC 27513

Work Phone: (919)677-8000

Email: Pat.Bostic@sas.com, Dave.Crow@sas.com, Clarke.Thacher@sas.com, John.Roth@sas.com, Arthur.Hunt@sas.com, Phillip.Smith@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.