

## The Platform for SAS<sup>®</sup> Business Analytics as a Centrally Managed Service

Joe Zilka, SAS Institute, Inc., Copley, OH  
Greg Henderson, SAS Institute Inc., Cary, NC

### ABSTRACT

Organizations that employ the Software as a Service model are looking to reduce the overall costs in supporting their “customers.” In the government sector, an Information Technology organization may be responsible for providing data integration, business intelligence, or analytic applications to several consuming agencies. A standardized set of hardware and software configurations as well as a set of best practices allows an organization to realize the benefits of centrally managed Software as a Service.

The popularity of logical partition and container capabilities provided by hardware and operating system vendors allows larger servers to be divided into many smaller “logical” servers providing consistency and ease of maintenance across the deployed software and applications. Sharing software binaries across the logical servers allows for enterprise-wide management of hot fixes and upgrades.

We will examine some of the advantages and benefits of deploying the platform for SAS<sup>®</sup> Business Analytics in a centrally managed environment. A discussion of some best practices and past experiences will also be presented.

### INTRODUCTION

Information Technology (IT) departments in all organizations are always looking for ways to reduce costs in the services they provide to their customers, especially in periods of reduced revenues and budget cuts in rough economic times. These customers are typically departments and agencies in both the private and public sector environments with which IT typically has a Service Level Agreement (SLA) to provide these services. Consistent use of hardware and software platforms provides economies of scale reducing costs for deployment and maintenance in providing a solution and service to the customer. The platform for SAS Business Analytics can be provided as a service to consuming departments and agencies while being managed from a central, IT-maintained environment.

### SOFTWARE DEPLOYMENT MODEL

The concept of Software as a Service has gained popularity in recent years as a computing option to reduce costs associated with a particular application. Software as a Service can be defined as a software deployment model where a service provider enables access to an application or set of tools that the provider supports, manages, and maintains in the provider’s computing environment. Many of the definitions you encounter also stipulate that access to the provided service is accessed only through Web-based interfaces such as a browser.

The advantages to this approach include but are not limited to the following:

- the ability to configure and deploy only the tools or modules required by the customer
- the ability to reuse binaries
- simplification of patch management tasks, enabling customers to quickly use the latest and greatest software versions
- allows multiple organizations to use a single set of the software binaries while each have their own data, metadata, and security models

A traditional software deployment entails installing, configuring, and maintaining application software on several hardware instances or containers. While allowing each instance to have its own software install and perhaps providing a segregated hardware and software environments, the management and maintenance of several instances can become a time-consuming task for IT shops that support many agencies or departments (tenants) using a wide range of capabilities of the platform for SAS Business Analytics.

## **CONFIGURING THE PLATFORM FOR SAS BUSINESS ANALYTICS IN A CENTRAL SERVICE MODEL**

The single most important requirement in centralizing the management of SAS software for use by multiple tenants is that of a shared file system for the install location of the SAS server tier software. Most hardware vendors with their own flavor of UNIX support some sort of file system sharing across their hardware partitioning technology. While the server side software will use a shared install, SAS Windows clients will have their own install for certain thin-client tools.

Creating instances of the platform for SAS Business Analytics requires an installation phase and a configuration phase. Installation is essentially copying the required applications to their desired location. Configuration would occur for each instance for a specific tenant and specific environment (development, test, or production).

### **HARDWARE PARTITIONING**

Hardware vendors such as IBM and Sun have technologies for their hardware offerings that allow a large server consisting of say, 144 cores, to be split up into many smaller servers with core counts ranging from single digits to tens of cores. IBM's Workload Partitioning and Sun's Container Model provide flexibility in how processors, memory, and bandwidth are allocated across the partitions or containers.

IBM's Workload Partitions (WPARs) slice a single operating system instance into multiple instances, which creates the logical partitions based on a single AIX OS. Each of the WPARs is its own isolated AIX environment preventing interference from another WPAR providing a secure and available environment from which applications can be hosted. Resources such as CPU and memory can be allocated and adjusted dynamically amongst the WPARs allowing flexibility in resource management between development, test, or production environments and tenant bases. Application WPARs can be quickly brought online with the applications installed on the global AIX environment thus reducing the time and IT resources required to make it available for use (Gibson 2007).

Very similar to WPARs, Sun Solaris Containers, or Zones, are a type of operating system-level virtualization technology that allows resource allocation and separation of environments providing fault isolation. A container is merely a zone that utilizes the resource management facility. These containers have their own node name, storage, and virtual network interface and do not require dedicated processor or memory; they can share portions of each. Separate zones can be created for different tenants and environments.

As mentioned previously, a file system upon which the SAS server tier software will be installed should be shared globally among the separate partitions. This allows a single set of binaries to be laid down and maintained. All partitions will be required to use the single set of executables. Large-scale storage environments typically involve a Storage Area Network (SAN) and use that technology to present file systems to the partitions. SANs provide a manageable way to organize and administer storage when servicing multiple SAS platform environments (development, test, or production) in a multiple tenant scenario.

### **SOFTWARE INSTALLATION**

Whether the file system dedicated for the application software is surfaced initially in a global partition or zone, or is made available to a local instance, the initial process of installing the SAS software is similar to a traditional server environment. Typically a SAS Software Depot, which is essentially a copy of the delivered software media, is created on the system. This depot expedites the installation process and would usually serve as the location from which all subsequent installs originate, but in this scenario the installation process takes place once. It is still recommended that a depot be created as access to a CD-ROM drive on the server is not always available during the installation process. The depot concept also helps to manage the software installation as service packs or new releases become available from SAS. In a multi-tier SAS deployment such as UNIX server housing the SAS application servers and middle tiers with a Windows client tier necessary for SAS clients such as SAS<sup>®</sup> Data Integration Studio, SAS<sup>®</sup> Information Map Studio, or SAS<sup>®</sup> Add-in for Microsoft Office, separate software depots will be necessary. The shared file system on the UNIX side will provide the necessary depot for the shared install. On the Windows desktops, the required SAS thin-client will be installed on the desktop for some user roles and a separate Windows shared location (i.e., network share) can be used as a depot from which these installs can be run. Web browser-based SAS clients require no software to be installed on the desktop.

The SAS Software Depot creation process along with the installation and configuration processes are thoroughly described in the SAS<sup>®</sup> 9.1.3 Intelligence Platform: Installation Guide (SAS Institute Inc. 2006). The SAS<sup>®</sup> Software Navigator available in the SAS Installation Kit should be used to create the requisite number of depots as described in the previous paragraph. If you receive a SAS Installation Kit for an updated version of the software, it is highly recommended that you create a separate depot for each of these versions. This will ensure the proper versions of all components are current. Once the depot has been created for the necessary systems, the installation of the software can begin.

The SAS® 9.1.3 Intelligence Platform: Installation Guide highly recommends that any required third-party products be installed prior to the SAS software installation. Third-party products that support the platform include Java Development Kit, servlet container or J2EE application server, WebDAV server, and scheduling software. Appropriate planning must be done to decide how these components will be installed in support of the SAS platform and the tenants and environments they serve. For example, a servlet container may be installed on each partition, which may in turn support a single tenant and one of the development, test, or production environments. This configuration will minimize the impact in the event of a servlet container crash and provide fault isolation. However, certain J2EE application servers provide the necessary isolation of server contexts, and it may be a feasible option to have a single environment instance (development, test, or production) supporting multiple tenants. A sample architecture of this configuration is discussed in a later section.

Much like the SAS Software Depot creation, the process of actually installing the SAS software to the desired file system is accomplished using the SAS Software Navigator. The installation process generally has two phases. The first step is selecting the software that is to be installed. The second step installs or prompts you to install the selected software. The list of software to install from which you select in the SAS Software Navigator wizard is determined by the PLAN.XML file that is created for your environment. In this case where we would be providing software to tenants, a plan file with all of the available server-side software for that operating system should be utilized. This ensures that the necessary components are made available to each tenant regardless of the components they decide to utilize and have deployed. Once all of the desired software components are selected, be sure to de-select the SAS Configuration Wizard as we do not want to configure any of the SAS software at this time. Configuration will occur on each tenant's partition specific to the tier and development, test, or production environments desired. The installation process should be repeated for any other operating systems used in a multi-tier, heterogeneous architecture. If the middle-tier servers are the same flavor of UNIX as the SAS application server, then the initial install would include the necessary middle-tier components. If a Windows middle tier is desired, then an install process would, of course, be repeated for that platform.

We have mentioned earlier the development, test, or production environments that may be deployed to support the requirements of a tenant's implementation of the SAS platform. It is important to note that these environments are for the purpose of development, test, or production of the end products of the platform for SAS Business Analytics. Tenants would most likely use a production set of binaries to support their platform development, test, or production deployments. Most IT shops would prefer to have a "test" location for the binaries in order to validate the application of updates and hot fixes. This can certainly be accommodated with a secondary installation location for the sole purpose of IT's testing of the software.

## **SOFTWARE CONFIGURATION**

Configuring the SAS software includes specifying the necessary information for the appropriate tier in order for the platform to be operational. This assumes that the necessary planning has been done to determine where the SAS® Metadata Server, SAS® application servers (SAS Workspace Server, SAS® Stored Process Server, and SAS® OLAP Server) and SAS middle-tier servers will reside. A typical four-tier deployment places the SAS Metadata Server on one host, the SAS application servers on a second, the SAS middle-tier servers on a third, and SAS clients on the last set of hosts. Decisions as to how third-party software will be managed will have an impact on how the configuration proceeds.

The SAS Configuration Wizard is used to gather the required information and configure the required components and is driven once again by the information in the PLAN.XML file. Therefore, a new PLAN.XML file will be created that reflects the product mix that each particular tenant requires. The appropriate planning is imperative to creating an environment with the proper components necessary for a successfully deployed and productive environment. In our example of a four-tier environment, the same plan file will be used to configure every tier and may also be used to configure any additional environments (development, test, or production) required by the tenant.

An important part of the input information gathered by the SAS Configuration Wizard in the centrally managed model is the location of the configuration directory. Each server partition will have its own configuration directory specific to its deployment. Deployment-specific information such as the Foundation metadata repository, logical server start-up scripts, log files and Web application files are all created by the SAS Configuration Wizard in the specified configuration directory. Therefore, each server partition should have its own private file system on which the configuration directory should be created. Moreover, a naming convention that allows for repeatable use of the configuration directory name and file system makes for easier maintenance of the configuration. For example, reviewing the metadata log for a particular deployment is simplified if a consistent, partition-specific location is specified for the configuration directory.

The SAS Configuration Wizard can be used to create multiple environments for a single tenant supporting the development, test, or production requirement. The flexibility of the platform for SAS Business Analytics allows these environments to coexist on a single server instance segregated by TCP/IP port assignment, or distributed across

additional servers. Distributing the environments across servers minimizes the possibilities of one environment impacting another and maximizes fault isolation.

## **ADDITIONAL CONSIDERATIONS**

As mentioned previously, dynamic resource allocation in the hardware environment can create a situation where hardware resources are utilized more efficiently and effectively. A set of servers can be controlled by a resource manager whereby CPU, memory, and bandwidth can be shared among those servers. The servers that are part of a resource manager group can all be of a particular tier of the platform (i.e., SAS Metadata Server, SAS application server, etc.) meaning that multiple tenants can have their metadata servers allocated a certain minimum resource set while utilizing a portion of any shared resources when conditions require.

Another consideration in deploying SAS as a centrally managed service is authentication. Depending on the end-user functionality, most users will be authenticated by some mechanism and authorizations determined thereafter. These authentication mechanisms are typically either host authentication or an alternative provider such as Lightweight Directory Access Protocol (LDAP) or Microsoft Active Directory. By configuring the SAS Metadata Server to use one of these alternative mechanisms, it negates the need to create individual user IDs on each operating system for the purposes of authentication. A user ID, either individual or group, will be required however to use a SAS Workspace Server. An operating system user ID assigned to a metadata user group reduces the number of user IDs needed and, in turn, the effort required to bring an environment online.

## **ADVANTAGES OF SAS AS A CENTRALLY MANAGED SERVICE**

The most obvious advantage of the centrally managed service model is that of manageability of the software once deployed. An IT organization can also reduce the time required to generate an environment that is ready for use by the requesting tenant. Another advantage is isolated environments for separate tenant entities as well as development, test, and production environments when those environments are configured individually.

The ability to run multiple SAS platform server configurations from a single install instance reduces the time and resources required to manage the current software deployment in comparison to a traditional model where the SAS platform software would be installed on every server required to support an IT shop's customer base. This is especially true when an IT organization must service several agencies. SAS does release Service Packs as part of a software update and enhancement strategy and also releases hot fixes to address specific system maintenance issues. The central software location for the SAS application server systems reduces the level of effort required to update systems when there is a requirement to update the application with a service pack or hot fix. Instead of performing the update three, four, or eight times, a single update can suffice.

Combining a robust hardware partitioning and resource management solution, SAS as a centrally managed service allows for rapid deployment of an IT customer's SAS environment. With the proper operating system and hardware configuration, SAS Metadata Server and application server environments can be brought online relatively quickly. Even though the configuration step is required for the application server, middle tier and installation and configuration is required for any thin-clients, the shared application server software can be quickly made available to the new environments. The time savings allows staff to concentrate on other details requiring attention such as the determining which capabilities are required and which tools satisfy the requirement.

## **A REAL-WORLD EXAMPLE OF THE SAS IN A CENTRAL SERVICE MODEL**

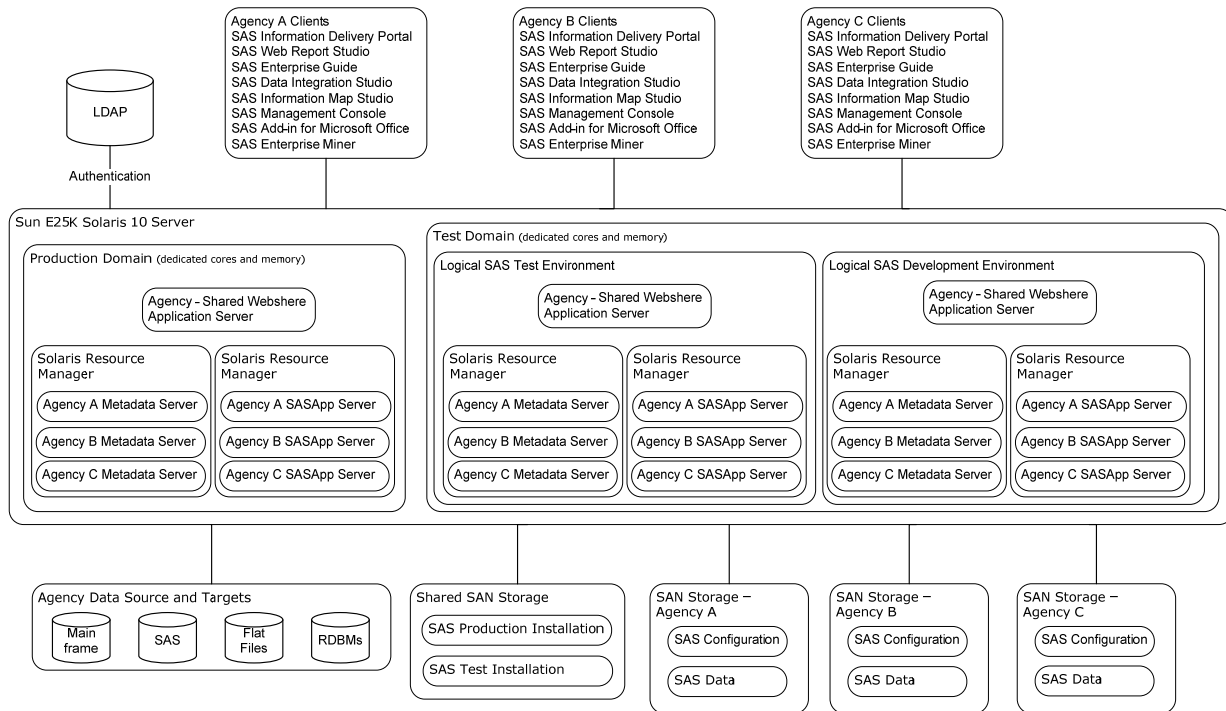
To obtain efficiencies and cost saving, the centralized IT department of a large state government deployed the SAS Intelligence Platform on a Sun Fire E25K server to service multiple disparate agency data integration, business intelligence and analytical applications. The goals of the IT department in delivering SAS in this manner was to provide the necessary autonomy to the agencies they supported, while at the same time making the most efficient use of resources and ensuring high service levels were maintained.

In order to accomplish these objectives, the SAS Intelligence Platform was installed on a single shared file system. Separate install images of the software binaries were created for development, test, and production, but these are then shared by all agencies. This configuration allows core SAS software patches and updates to be vetted through a testing process before being applied in the production environment.

The server was physically partitioned into Development/Test and Production domains, each running Solaris 10. Within each domain, logical containers were created to support the individual agency applications. By configuring things in this manner, the Solaris 10 Resource Manager can be leveraged to set minimum resource assignments to each container to meet SLAs, but can then dynamically allocate unused resources from one agency container to another depending on current utilization and workloads for each container.

The middle-tier J2EE server is WebSphere 6, hosted on the same E25K server. WebSphere Managed Servers are utilized to provide a separate Web container for each agency's Web applications. This allows customized branding and other customizations to the core SAS Web application to occur independently for each agency. Xythos WebDAV server is installed on the same container as WebSphere.

Authentication is provided by an external LDAP server, which is integrated and synchronized with a statewide enterprise identity management system. To facilitate ease of administration of host-based user accounts, most metadata identities do not contain individual logons. Instead, identities are arranged into groups, and logons are maintained at the group level. This approach minimizes the need for synchronizing passwords on the user account level, but does have the disadvantage of not being able to track each individual user's resource consumption at the OS level.



**Figure 1. Case Study Architecture**

For each new “customer” or agency that the IT department brings on board to the service, a standard process is repeated to create an environment for that customer. First, the appropriate containers are created on the server. Standard naming conventions are used to facilitate ease of deployment and administration. Next, the PLAN.XML file is created and the SAS Configuration Wizard is run. The result of this process is that a configuration directory is created specific to this agency application. This configuration directory contains all of the necessary server startup scripts (specifying the correct host names, IP addresses, and ports for the agency's containers), configuration files, and metadata. Finally, the SAS Web applications are deployed to the Web application server. Since the environment contains development, test and production environments, the configuration process above is repeated for each environment.

In this particular case study, the IT department's role is to provide the necessary SAS platform infrastructure for the agencies to leverage. It is the agency's responsibility to then utilize the tools and resources provided by environment to create “content” such as data integration, business intelligence, or analytical applications. As mentioned previously, separate development, test, and production environments are provided to ensure that application changes can be appropriate vetted prior to impacting a production system. To facilitate this process, standard process and scripts were deployed to allow for migration of content from one environment to another. The scripts provide for copying of certain data or content items, such as source code files, from one environment to another. Both full and partial metadata promotion utilities are utilized to move metadata across the different environments.

In a typical scenario for a new customer, the development, test, and production environments would be provisioned and configured by the IT department. The agency would then utilize tools such as SAS Data Integration Studio, SAS Information Map Studio, SAS® Web Report Studio, and SAS® Information Delivery Portal to create content and applications in the development environment. The content from the development environment is then “promoted” to the test environment using full metadata promotion and the custom scripts developed for promoting other associated content items. If any changes result from integration testing, those changes are made in the development environment and then promoted again to test. When the final application is ready for production deployment, a similar promotion occurs from either the development or the test environment to the production environment.

Because the SAS Intelligence Platform provides the ability for end users to create their own content directly in the production environment, care must be taken to ensure that these elements are not destroyed on subsequent dev -> test -> prod promotions. To facilitate this, after the initial deployment into production, all subsequent promotions should utilize partial promotion of only the content items that have been updated.

## CONCLUSION

Creating the platform for SAS Business Analytics as a centrally managed service requires a carefully planned out architecture for the hardware and software perspectives. However, over time this perpetuates a robust, manageable and efficient set of applications that are easier to manage and faster to deploy yet they remain fault tolerant in relation to each other. This type of architecture is better suited to an enterprise-level deployment in that economies of scale provide much of the cost savings realized. This approach was employed by an organization and has found it an effective method of deploying the platform for SAS Business Analytics to multiple tenants requiring development and test environments on a server, while delivering a production environment on another server.

## REFERENCES

- Gibson, Chris. “WPAR Power.” *IBM Systems Magazine*, December 2007. Available at <http://www.ibmssystemsmag.com/aix/december07/coverstory/18606p1.aspx>.
- SAS Institute Inc. 2006. *SAS® 9.1.3 Intelligence Platform: Installation Guide, Fifth Edition*. Cary, NC: SAS Institute Inc.

## RECOMMENDED READING

SAS® 9.1.3 Intelligence Platform: Installation Guide, Fifth Edition,  
<http://support.sas.com/documentation/configuration/biig.pdf>

SAS® 9.1.3 Intelligence Platform: Security Administration Guide, Second Edition,  
<http://support.sas.com/documentation/configuration/bisecag.pdf>

## CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the author:

Joe Zilka  
SAS Institute, Inc.  
5099 Fairington Ave  
Copley, OH 44321  
Work Phone: 330.598.1421  
E-mail: [Joe.Zilka@sas.com](mailto:Joe.Zilka@sas.com)  
Web: <http://www.sas.com>

Greg Henderson  
SAS Institute, Inc.  
Cary, NC 27514  
E-mail: [Greg.Henderson@sas.com](mailto:Greg.Henderson@sas.com)  
Web: <http://www.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.