Paper SAS298-2014

SAS® Visual Analytics for the Three Cs: Cloud, Consumerization, and Collaboration

Christopher Redpath and Nicholas Eayrs, SAS Institute Inc., Cary, NC

ABSTRACT

SAS® Visual Analytics delivers the power of approachable in-memory analytics in an intuitive web interface. The scalable technology behind SAS Visual Analytics should not just benefit the analyst or data scientist in your organization, but, indeed, everyone regardless of their analytical backgrounds. This paper outlines a framework for the creation of a cloud deployment of SAS Visual Analytics using the SAS® 9.4 platform. Based on proven best practices and existing customer implementations, the paper focuses on architecture, processes, and design for reliability and scalable multi-tenancy. The framework allows your organization to move away from the departmental view of the world and offer analytical capabilities for consumerization and collaboration across the enterprise.

INTRODUCTION

The current economic climate dictates that you need to be swift and agile in your response to ever-increasing customer expectations and competition. This requires you to maximize every competitive advantage that you can. Your data is just that (i.e., a free-to-use resource and key differentiator promising unrivaled insight into your business and its improvement). Realizing the true potential of your data is often challenging because of the lack of expertise or capability to access, explore, and derive insight from the data itself, as well as to share insight for multi-disciplinary collaboration. SAS Visual Analytics addresses this very challenge by empowering you with self-service advanced analytics capabilities geared for driving collaboration and information sharing via a flexible deployment model that leverages the cloud.

CLOUD COMPUTING CHARACTERISTICS, SERVICE, AND DEPLOYMENT MODELS

For the purpose of this paper, here is the National Institute of Standards and Technology (NIST) definition of cloud computing:

"Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction." [1]

CLOUD COMPUTING CHARACTERISTICS

The NIST definition defines five essential characteristics of the cloud computing model:

1. **On-demand self-service—**consumers are able to provision software effectively and efficiently.
2. **Broad network access—**consumers can access software capabilities over the network, leveraging standards-based connectivity that supports a diverse set of clients (i.e., thin or thick (e.g., tablets, mobile phones, laptops, or workstations)).
3. **Resource pooling—**consumers share resources that serve other consumers in a secure model that allows for economies of scale.
4. **Rapid elasticity—**consumers are able to scale resources as needed.
5. **Measured service—**consumers can leverage metered software capabilities (i.e., capacity-based, user-based, and value-based pricing).

CLOUD COMPUTING SERVICE MODELS

The NIST definition defines three cloud computing service models:

1. **SaaS—**consumers are able to access software capabilities over the network, typically leveraging nothing more than a thin client (i.e., a web browser). The consumer does not manage or control the underlying cloud infrastructure or software configuration, but is able to customize user-specific settings within the software.
2. **PaaS—**consumers are able to deploy self-developed or acquired software using APIs, services, and tools supported by the cloud service provider (CSP). The consumer does not manage or control the underlying cloud infrastructure, but is typically able to customize software deployment and configuration settings.
3. **IaaS—**consumers are able to provision processing, storage, networks, and other fundamental computing resources for the development and deployment of software. The consumer does not manage or control the

underlying cloud infrastructure, but is typically able to customize the operating system, deployed software, and the network configuration.

CLOUD COMPUTING DEPLOYMENT MODELS

The NIST definition defines four cloud computing deployment models:

1. **Private cloud—**cloud infrastructure provisioned for the exclusive use of a single organization encompassing multiple consumers. It can be owned, managed, and operated by the organization, a third party, or some combination, and it can reside on or off premises.
2. **Community cloud—**cloud infrastructure provisioned for the exclusive use of a community of consumers from organizations that have shared concerns (e.g., regulatory compliance, security policies, etc.). It can be owned, managed, and operated by one or more organizations within the community, a third party, or some combination, and it can reside on or off premises.
3. **Public cloud—**cloud infrastructure provisioned for use by the public. It can be owned, managed, and operated by a business, academic or government organization, or some combination, and it resides on the premises of the CSP.
4. **Hybrid cloud—**the coalescence of the cloud infrastructure from two or more distinct cloud deployment models that remain unique entities for the purpose of data and application portability (e.g., cloud bursting).

THE SAS CLOUD

The SAS Cloud is a hosted, private cloud that provides rapid access to SAS software and an easy way to manage your SAS environments. The SAS Cloud characteristics align with the NIST definition. In addition, it supports two of the three NIST cloud computing service models, SaaS and PaaS. The IaaS cloud computing service model is currently not supported by SAS because of our continued focus on being a software provider as opposed to an infrastructure provider.

Within the SAS Cloud, your organization is licensed to use one or more dedicated SAS environments. A SAS environment is a complete, self-contained, fully functioning deployment of SAS applications. Environments in the SAS Cloud are deployed using virtualization technology, which makes the environments easy to create and manage.

A typical environment includes the following:

- data storage that can be populated with your organization’s data
- virtualized software, servers, and services that perform various types of processing
- virtualized client applications through which users can access SAS functionality
- administrative interfaces for uploading data, organizing SAS content, and managing permissions

SAS CLOUD TOOLS AND TECHNOLOGIES

The SAS Cloud relies on virtualization technology to ease the deployment and maintenance of environments. The essential tools and technologies includes the following:

- **SAS Virtual Applications (vApps)—**a collection of SAS software, third-party components, and configuration settings that are packaged to provide a ready-to-run SAS application environment. The vApp is deployed in a virtualization environment and external connectivity is established accordingly (e.g., security, monitoring, and data storage). SAS Visual Analytics is the first SAS solution to be offered as a vApp.

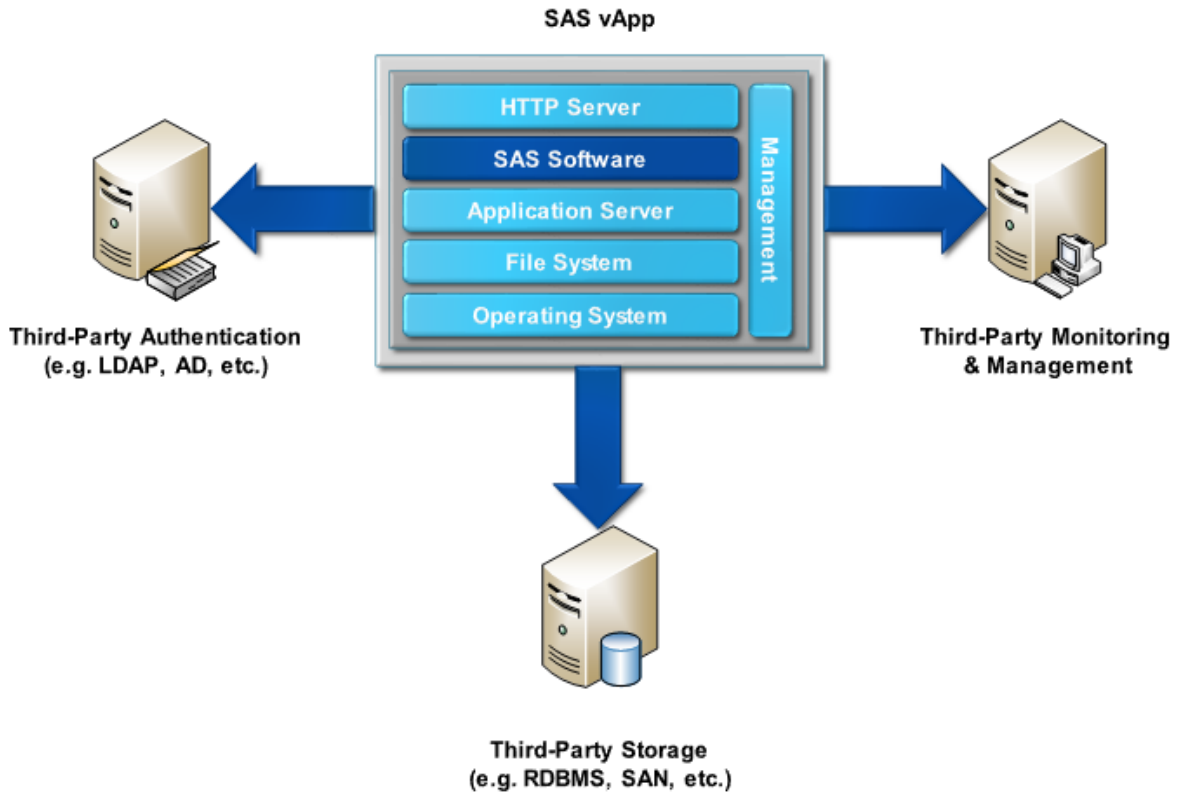

Figure 1. Components of a SAS vApp

- **SAS App Central—**a role-based web portal for managing your SAS Cloud account, applications, users, and environments and for uploading your data. This is the default landing page for SAS Visual Analytics for SAS Cloud.

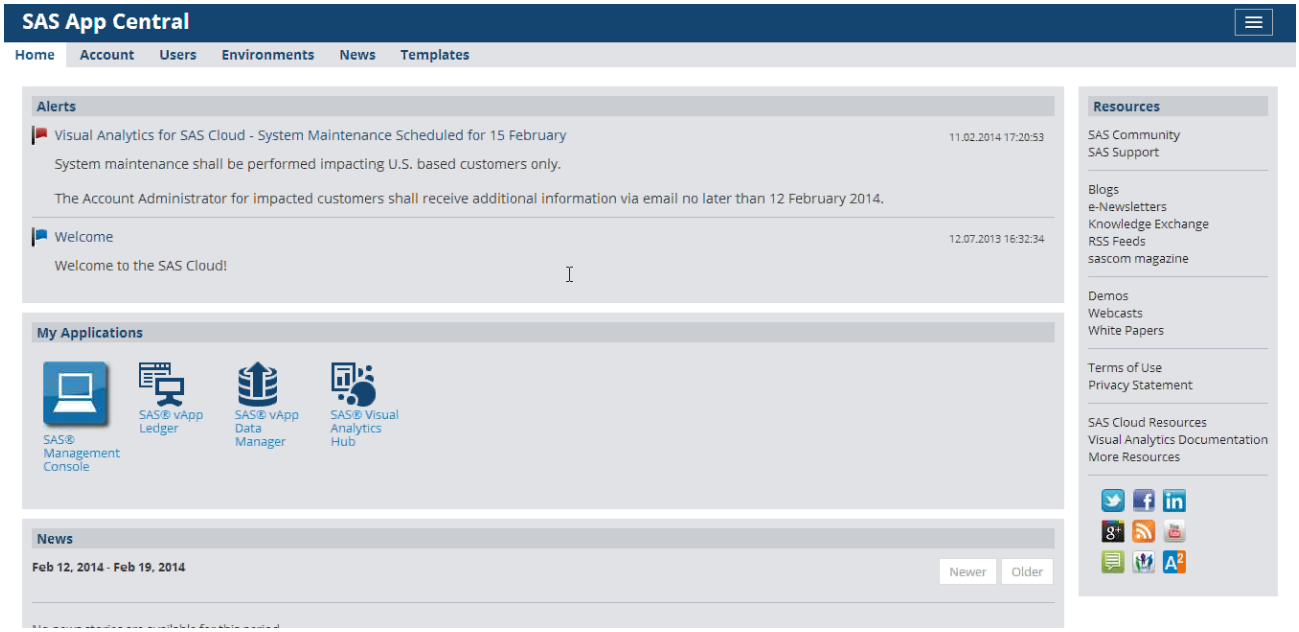

Figure 2. SAS App Central Initial Screen

- **SAS App Engine—**a distributed technology that enables the assembly, delivery, and maintenance of SAS vApps through SAS App Central. This facilitates the automated provisioning of SAS Visual Analytics for SAS Cloud as well as providing push notifications of the latest SAS software updates for you to selectively choose.

SAS VISUAL ANALYTICS CLOUD DEPLOYMENT MODELS

There are several deployment options for SAS Visual Analytics. Whether you prefer to deploy the solution on-site using your own hardware or via private clouds, public clouds, or the SAS Cloud environment, we have a deployment model to fit your business.

SAS VISUAL ANALYTICS FOR SAS CLOUD

SAS Visual Analytics for SAS Cloud is the perfect choice if you do not have the time or resources to support a traditional, on-site implementation. It provides an affordable license model using predefined, non-distributed SAS Visual Analytics environment sizes, also known as team sizes, within the SAS Cloud.

The following is included with each license:

- Cloud infrastructure (software and hardware)
- Express Support Service, which consists of loading up to three data tables
- Standard SAS Technical Support services
- Access to online documentation
- Access to SAS Visual Analytics Video Library

The various team sizes are summarized in the following table:

	Small	Medium	Large
Total Named Users (Unique user IDs)	5	10	15
Data Capacity (GB) (Data stored on disk)	32	64	128
Usable Storage (GB) (Data that can be loaded to memory)	16	32	64
Maximum Table Size (GB) (Maximum size of single table)	4	6	8

Table 1. Team Sizes for SAS Visual Analytics for SAS Cloud

With this deployment model, you simply need to liaise with your local SAS representative to procure a SAS Visual Analytics for SAS Cloud environment and accept the associated electronic click-through legal agreement. Your environment will be automatically provisioned for use.

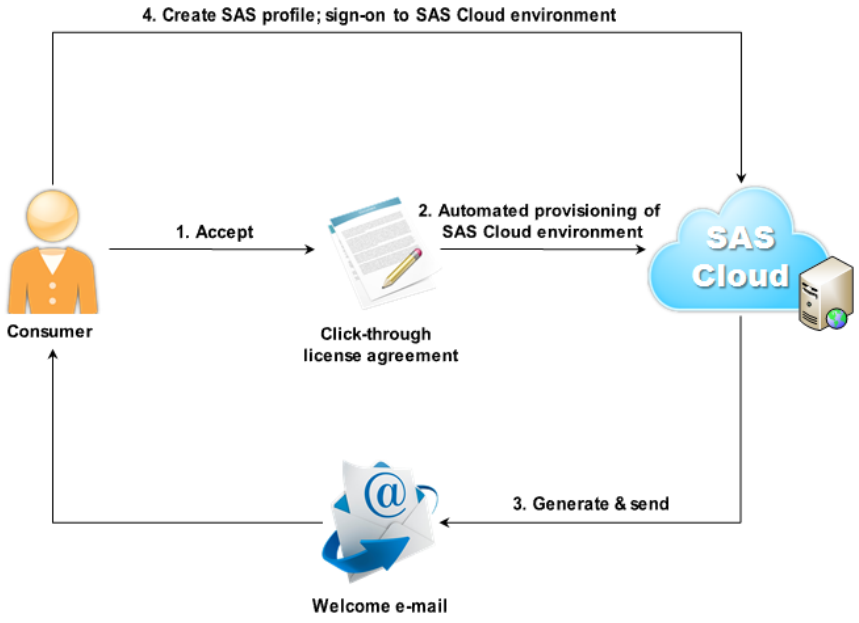

Figure 3. High Level Process for Provisioning SAS Visual Analytics in the SAS Cloud

All that is needed from this point on is Internet connectivity, a supported web browser [2], your data, and a business problem to explore.

Typical tasks to configure or customize your SAS Visual Analytics for SAS Cloud are as follows:

- Add additional users or administrators up to the maximum specified by your team size.
- Assign users to predefined SAS Visual Analytics for SAS Cloud roles as follows:

- **Data administrator—**Data administrators have the ability to load data into the SAS Visual Analytics environment using the SAS vApp Data Manager application.
- **SAS administrator—**SAS administrators have the ability to access administrative applications, such as SAS Management Console and SAS vApp Ledger.
- **SAS user—**SAS users have the ability to access SAS applications as regular users.
- **Account administrator—**Account administrators have the ability to maintain the customer's account, including adding users, providing users access to applications, and assigning users to roles. The account administrator is also able to manage environments.

- Prepare and load data using either the SAS Visual Data Builder or the SAS vApp Data Manager application. For more information, see the *SAS Visual Analytics for the SAS Cloud: Quick-Start Guide*. [3]
- Configure SAS Mobile BI for iPad or Android-based tablets to access your SAS Visual Analytics reports. For more information, see the *SAS Visual Analytics for the SAS Cloud: Quick-Start Guide*.

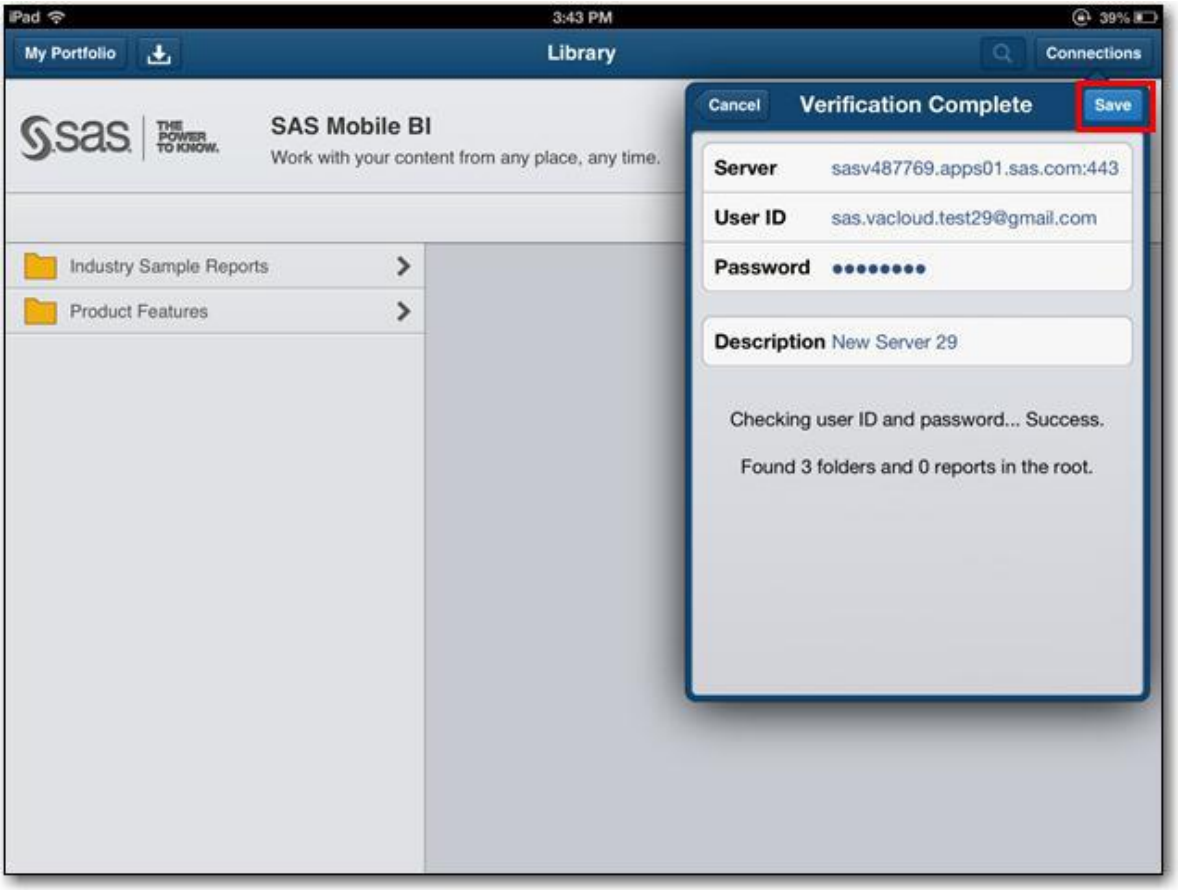

Figure 4. SAS Mobile BI Connection to the SAS Cloud

SAS Visual Analytics for SAS Cloud is available in the United States, Europe, the Middle East, and Africa (EMEA).

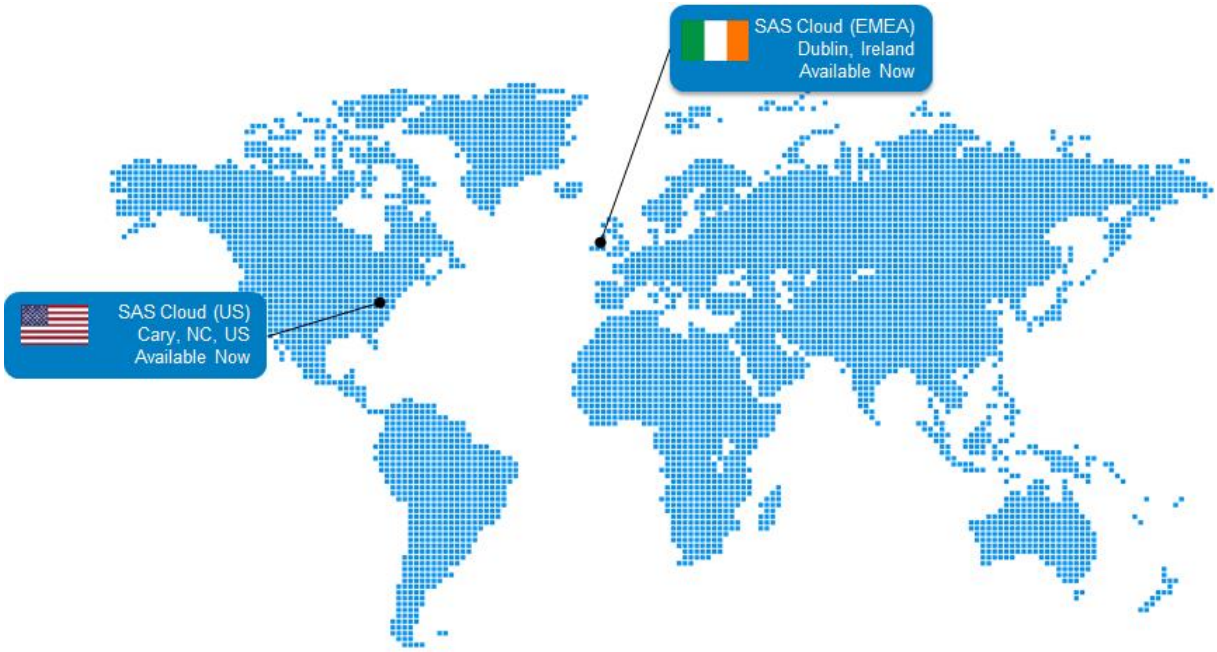

Figure 5. SAS Cloud Locations

SAS VISUAL ANALYTICS FOR PUBLIC CLOUD

SAS Visual Analytics for Public Cloud currently supports SAS Visual Analytics distributed and non-distributed deployments to Amazon Web Services (AWS) only (i.e., Amazon Elastic Compute Cloud (EC2) and Amazon Virtual Private Cloud (VPC). Although EC2-Classic is supported as a deployment model, it is recommended that customers use VPC, which provides increased isolation, control, customization, and security through a virtual private cloud environment within Amazon's public cloud infrastructure.

Different from the SAS Cloud model, with this approach, SAS Visual Analytics is installed on Amazon's public cloud infrastructure using the same process as the on-premises deployment of SAS. However, there are configuration and integration differences between the traditional on-premises deployment of SAS and Amazon's deployment of SAS, which are highlighted below.

With the public cloud deployment model, it is strongly recommended that you develop a SAS Visual Analytics for Public Cloud architecture customized for your business with support from SAS Professional Services.

Once the target architecture for SAS Visual Analytics is defined, you need to complete the following steps from your AWS Management Console:

Figure 6. Steps for Amazon Deployment of SAS Visual Analytics

1. **Choose AMI—**choose an Amazon Machine Image (AMI) with a virtualization type equal to the hardware virtual machine (HVM). This is the only supported option for cr1.8xlarge instances. (See step 2 for more information.) In addition, you must select a supported operating system for your SAS Visual Analytics distributed or non-distributed deployment as follows:
    - SAS Visual Analytics Distributed:
        - Red Hat Enterprise Linux 6
        - Suse Linux Enterprise Server 11
    - SAS Visual Analytics Non-Distributed:
        - Red Hat Enterprise Linux 6
        - Suse Linux Enterprise Server 11
        - Windows Server 2008 R2 (Enterprise or Datacenter)
        - Windows Server 2012 (Standard or Datacenter)

2. **Choose Instance Type—**For now, the only supported instance type for SAS Visual Analytics for Public Cloud is cr1.8xlarge.

Size	ECUs	vCPUs	Memory (GiB)	Instance Storage (GiB)	EBS-Optimized Available	Network Performance
cr1.8xlarge	88	32	244	2 x 120 (SSD)		10 GB

Table 2. Amazon Instance Specifications

Note: The cr1.8xlarge instance type is currently available only within the following regions (as of this writing):

- US-East/US Standard (Virginia)
- US-West-2 (Oregon)
- EU (Ireland)
- Asia Pacific (Japan)
- Asia Pacific (Tokyo)

3. **Configure Instance Details—**configure the instances to meet your business needs as follows:
    - Select the **Number of Instances** that you intend to launch for either a SAS Visual Analytics distributed or non-distributed deployment (i.e., you can launch more than one instance at a time).
    - Select the **Purchasing Option** for instances (e.g., spot or on-demand).
    - Select your preferred **Network** to launch instances either into EC2-Classic or VPC.
    - Select your preferred **Availability Zone** to specify a distinct location within a region.
    - Select a **Placement Group** to specify a logical grouping for your instances with low-latency, full bisection, 10-GBps connectivity between them.
    - Select an **IAM Role**, allowing you to reuse your existing Amazon credentials for the launch of your instances.
    - Select your preferred **Shutdown Behavior**. Instances can either be stopped or terminated at shutdown.
    - **Enable termination protection** to protect against accidental termination.
    - Enable **Monitoring** to allow detailed monitoring via CloudWatch (additional charges apply).
4. **Add Storage—**enables you to attach additional Elastic Block Store (EBS) volumes and instance store volumes to your instance or to edit the settings of the root volume. If you are uncertain about sizing, contact your local SAS representative to arrange a SAS sizing exercise for your specific workload requirements.
5. **Tag Instance—**enables you to categorize your instances using key and value pairs. This can be an invaluable option for ease of management. For more information about the basics of tagging, restrictions, and constraints, see *Tagging Your Amazon EC2 Resources*. [4]
6. **Configure Security Group—**security groups need to be modified to allow traffic to flow between hosts and to ensure that you do not expose your environment unnecessarily to the Internet, which will pose security risks unless mitigation and consideration has been given to hardening your environment.
7. **Review Instance Launch—**review your instance launch details before launching and take note of any warnings that might include security or pricing specifics.

Once connectivity to the instances is validated using your Amazon credentials, you must complete all of the necessary prerequisite steps from the SAS Visual Analytics deployment and administration documentation. [5] If all prerequisites are met, then the deployment of SAS Visual Analytics can be completed using the instructions in the deployment and administration documentation.

Additional public cloud deployment model considerations are:

- Create a DNS Alias with a CNAME (Canonical Name) record in DNS for each host, pointing to the Elastic IP address, and use that as the fixed host name.
- Exposing your SAS Visual Analytics environment to the Internet can be risky unless mitigation and consideration has been given to hardening your environment. If in doubt, please contact your local SAS representative to arrange for support from SAS Professional Services.
- The time and cost to upload data need to be carefully considered in such a deployment.

For more details, consult your local SAS Professional Services before embarking on a SAS Visual Analytics for Public Cloud deployment.

SAS VISUAL ANALYTICS FOR PRIVATE CLOUD

SAS Visual Analytics for Private Cloud supports SAS Visual Analytics distributed and non-distributed deployments in virtualized environments using the VMware vCloud Suite, which provides an open and modular architecture for the development of public or private clouds.

The VMware vCloud Suite includes the following technologies:

- VMware vCloud Director and API – orchestrates the provisioning of virtual data centres
- VMware vSphere – provides the virtualisation platform for the building of your cloud infrastructure
- VMware vCloud Networking and Security – provides networking and security capabilities for your cloud
- VMware vCenter Operations Management Suite – provides insights into the performance, capacity and health of your cloud infrastructure
- VMware vCloud Application Director – allows you to build deployment blue prints for automated application deployment using a DevOps model
- VMware vCenter Orchestrator – allows you to design and automate complex workflows
- VMware vCloud Connector - links your public and private clouds so that you can easily manage them as a hybrid cloud

As you can see by the extensive virtualization capabilities available to you, you have full control over the resources within your private cloud and the provisioning at the most granular level possible. Therefore, careful consideration must be given to mitigate resource contention and to achieve optimal performance for SAS Visual Analytics.

SAS strongly recommends leveraging the VMware *vCloud Architecture Toolkit (vCAT) 3.1.2* [6] body of knowledge to develop an appropriate architecture for your specific needs.

Figure 7. VMware vCAT 5 Steps

Once your virtual machine instances are deployed and connectivity to the instances is validated using your supplied credentials, you must complete all of the necessary prerequisite steps from the SAS Visual Analytics deployment and administration documentation. If all prerequisites are met, then the deployment of SAS Visual Analytics can be completed using the instructions in the deployment and administration documentation.

Additional private cloud deployment considerations are:

- VMware vSphere configuration maximums must be considered.

The following configuration maximums are relevant to SAS Visual Analytics for vSphere 5.5:

Item	Maximum
Maximum number of virtual CPUs per virtual machine (Virtual SMP)	64
RAM per virtual machine	1TB

Table 3. Virtual Machine Maximums

Item	Maximum
Maximum number of virtual CPUs per host	4096
Maximum number of virtual machines per host	512
Virtual CPUs per core	32

Table 4. Compute Maximums

Item	Maximum
Volume size (common to all VMFS versions)	64TB
VMFS3 file size (default block size = 8 MB)	2TB minus 512 bytes
VMFS5 file size (default block size = 1 MB)	62TB

Table 5. Storage Maximums

Item	Maximum
RAM per host	4TB

Table 6. Memory Maximums

- In the BIOS:
    - Enable hardware-assisted CPU virtualization (i.e., either Intel (VT-x) or AMD-V).
    - Enable hyperthreading.
    - Enable Turbo Boost, if supported.
    - Disable C1E and other C-states in BIOS.
- Do not over-provision resources.
- Avoid thin provisioning of disks, default to eager zeroed.
- Enable jumbo frames for networking.
- Use VMXNET3 virtual network adapters.
- Set the ESXi power policy to high or maximum performance.
- Install the latest versions of VMware tools in the guest operating system.
- Properly size storage environments to accommodate worst-case scenarios (such as maximum required MB per second I/O throughput for a virtual host).
- Design the system to incorporate extra resources to overcome virtualization overhead.
- Use VMware monitoring tools to help diagnose performance issues at the virtual host level.

If you are using Red Hat Enterprise Linux as the guest operating system, perform the following steps:

- Set the Red Hat Enterprise Linux I/O elevator to deadline.
- I/O barriers can be safely disabled with enterprise storage whose cache controller RAM is battery backed up. (You can switch the I/O elevator to deadline using the tuned tool.)
- Enable the performance CPUfreq governor. [7]
- Disable transparent huge pages because large file I/O might perform better with this feature disabled. This can be done with the following command:

  ```
  # echo never > /sys/kernel/mm/redhat_transparent_hugepage/enabled
  ```

  Note: The tuned tool enables this feature by default for all profiles, so anytime a tuned profile is enabled, whether automatically at start up via an init.d service or manually at the command line, this feature will be re-enabled.

For more information, consult your local SAS Professional Services before embarking on a SAS Visual Analytics for Private Cloud deployment.

CONSUMERIZATION

Consumerization is the emergence of design patterns around the consumer, not the organization. However, because of trends such as cloud computing, *app*-ification, and mobility, it is having a profound impact on the organization. Organizations must embrace consumerization by empowering their users to be productive with personal devices or risk users seeking access to on-demand self-service software capabilities outside of the organization. These are often in the form of a personal cloud connected to a personal device, typically a tablet.

With SAS Visual Analytics, you are able to support the broadest range of consumer devices with SAS Mobile BI client support for iPad and Android tablets.

SAS Mobile BI client support includes:

- iOS
    - SAS Mobile BI for iPad is a free application for iOS available in the iTunes App Store.
    - Platform or OS: Apple iOS 6.0 and above
    - Devices: iPad 2, 3, 4, and Mini
- Android
    - SAS Mobile BI for Android is a free application for Android available from Google Play.
    - Platform/OS: Android 4.1 and above
    - Devices: Android tablets 10.1

COLLABORATION

ANALYTICS CENTER OF EXCELLENCE (COE)

To ensure the continued realization of value from your investment in SAS Visual Analytics, it is imperative that your focus is not only on the technology considerations, but, more importantly, on the people or consumers. A key enabler in achieving user adoption is to enable individuals to feel empowered and free to collaborate on a shared analytic vision supported by your executives. SAS can support you in establishing a business analytics center of excellence (CoE) to achieve this. [8]

MULTI-TENANCY IN A PRIVATE CLOUD—CUSTOMER SCENARIO

The following scenario is based on a multi-tenanted private cloud deployment of SAS Visual Analytics. This deployment is built in distributed mode with Hadoop as the co-located data provider. The processes adopted can be used as a guide to achieve a similar lightweight framework to handle multiple tenants. However, it is only one example of a possible framework to meet specific requirements.

At a high level, the main requirements are:

- A centrally IT-provisioned and managed private cloud platform that offers SAS Visual Analytics as a SaaS to the business. This includes SAS Mobile BI access from tablets.
- The platform has to be scalable, starting with a smaller cluster and growing to meet the business demand as more tenants come on board.
- Where required by the tenant, there needs to be end-to-end data segregation from other tenants.
- There is support for sharing content to promote collaboration between tenants, especially tenants belonging to the same parent business unit.
- There is support of an appropriate chargeback model to fairly bill-back tenants for usage and to drive the investment required to maintain and grow the platform.
- IT manages batch loading of data into the environment from various source systems.
- The business is able to interactively load its own data into SAS Visual Analytics.
- There is support for scripting to manage the bulk of the technical onboarding process for new tenants.
- There is support for a lightweight authentication framework with minimal maintenance.

MULTI-TENANCY

For the purpose of the scenario, we used the following definition of multi-tenancy provided by Forrester [9]:

“Multi-tenancy defines IT architectures that let multiple customers (tenants) share the same applications and/or compute resources with security, reliability, and consistent performance.”

In line with this definition, the SAS Visual Analytics platform provides shared application software and compute resources for use by the tenants.

Advantages of the multi-tenancy approach for SAS Visual Analytics include:

- A centrally managed platform under the supervision of a core nucleus of employees. It assures lower costs, a higher quality of SAS business services, and a better match between availability and resources.
- The average cost for the development of SAS Visual Analytics services provided by central IT is lower than individually deployed SAS Visual Analytics environments. Business units can get started more quickly because they do not need to build and fund an entire new platform. Costs are lowered by using centrally located and maintained infrastructure and software to reduce up-front setup and ongoing operating costs for each participating tenant.
- Higher utilization of the hardware results directly in lower infrastructure costs because hardware is never idle.
- Building a platform based on standard, modular, and reusable components provides the capability to leverage effectiveness and efficiency by sharing best practices in terms of concepts, applications, infrastructure, and services among all tenants.
- Users can benefit from the large capacity of a shared environment. For SAS Visual Analytics, this means that more data can be loaded into memory, larger data sets can be analyzed, and more computing resources are available to provide faster results. Having a single, large enterprise infrastructure shared among many tenants means shared benefits for all.
- Software upgrades and patches can be applied quicker and with less effort because there is only one instance of the software to update (as opposed to sharing infrastructure instead of sharing the application).

SAS LASR ANALYTIC SERVER AS THE COMPUTATIONAL PLATFORM

There are important differences between building a multi-tenanted environment for traditional SAS code execution (i.e., a SAS grid-based environment) and an environment for SAS Visual Analytics [10] [11]. The primary difference is user activity within SAS Visual Analytics web clients includes interacting with and launching data in the SAS LASR Analytic Server. This is contrary to the traditional SAS grid where the SAS Workspace Server is the computational layer for all user requests. Therefore, the shared computational platform for users in a SAS Visual Analytics environment is the SAS LASR Analytic Server cluster.

The standard flavor of the SAS LASR Analytic Server in a distributed environment supports the Hadoop Distributed File System (HDFS) for co-located storage of data with replication and failover capabilities. The smallest configuration of the SAS LASR Analytic Server in distributed form contains four nodes. The availability of local storage on the nodes is critical for storing large data sets in distributed form in HDFS.

The SAS LASR Analytic Server cluster consists of the following components:

- A LASR root node. The server that contains the LASR root node is also the Hadoop NameNode when Hadoop is your co-located data provider.
- LASR worker nodes (all of the computational nodes that will hold data). Each machine that contains a worker node is also an HDFS data node when Hadoop is your co-located data provider.

The topology for SAS Visual Analytics encompasses the client tier, middle tier, and server tier. These tiers interact with the SAS LASR Analytic Server to enable client access to data for performing analytic tasks. Any actions initiated by clients are channeled by the middle tier and the server tier before they are sent to the SAS LASR Analytic Server. In a distributed environment, the LASR root node receives the processing action request and sends it to the LASR worker nodes. Each LASR worker node processes the action in parallel on its portion of the data before sending the results back to the LASR root node for the final aggregation to bring all the worker node results together. The final result goes back to the client via the middle tier.

BUSINESS UNITS AND BUSINESS DATA AREAS

In this scenario, for the multi-tenanted platform, the tenants can be organized as belonging to a two-level hierarchy:

- Business unit: Within the scope of the scenario, this is defined as a top-level division within an organization (for example, Retail, Marketing, Finance etc.).
- Business data area: Each business unit (BU) can have a number of distinct areas of data that need to be physically segregated, possibly corresponding to functions within a BU. For example, Retail can have separate data areas for Retail Marketing and Retail Brand. These are referred to as business data areas

(BDAs). Each of these can have its own data access controls and SAS Visual Analytics content. In addition, BDAs can relate to segregated projects within the BU.

BUs and BDAs can also be virtual. They do not have to strictly fit and restrict users based on the organization structure. For example, virtual BUs and BDAs can help foster analysis on a common set of data between tenants in the same organization.

In terms of data and content access security, this can be enforced at both the BU level and at the BDA level. Therefore, a user would need to be an identified member of a BU and a BDA to be able to access data secured at the BDA level. Moving one level up, it is possible for non-restricted data within a BU to be generally available to all tenants at the BU level. Moreover, group-wide non-restricted data can be made available to all SAS Visual Analytics tenants if required. Finally, at all levels, row-level data restrictions could be enforced.

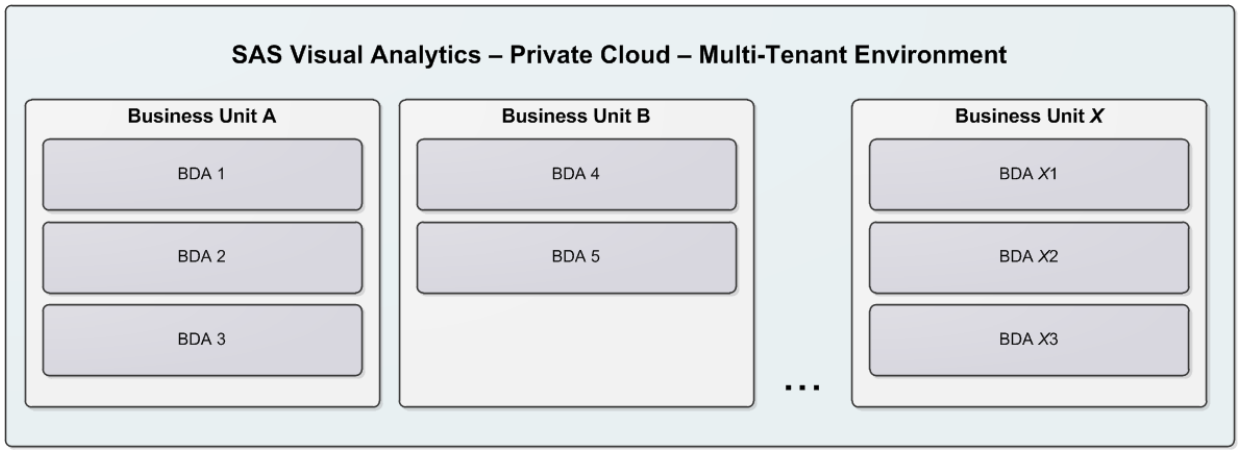

Figure 8. Business Units and Business Data Areas Hierarchy in a Multi-Tenant Environment

METADATA SECURITY AND FOLDER STRUCTURE

The SAS 9.4 platform provides a metadata-based authorization layer that supplements protection from the host environment and other systems. You can use this layer to manage access to almost any metadata object (for example, reports, folders, data definitions, explorations, jobs, stored processes, and server definitions).

SAS clients use a hierarchy of SAS folders to store and organize metadata. When you install SAS, a set of default SAS folders is created. The folders are arranged in a structure that segregates system information from business information, provides personal folders for individual users, and provides an area for shared data. Within this overall structure, a customized folder structure that meets specific information management, data sharing, and security requirements can be created.

The folder structure below is an example base template for the SAS Visual Analytics multi-tenanted environment, followed by a description of the various subcomponents of the structure. It represents a structure that allows for easy implementation of metadata security. It is a generic structure that helps script the process of onboarding a tenant and can be expanded or modified depending on future needs and to accommodate more projects.

The folder structure should be protected using a SAS metadata security model that uses access control templates (ACTs). An ACT is a reusable, named authorization pattern that can be applied to multiple resources. An ACT consists of a list of users and groups that indicates whether permissions are granted or denied for each user and group. [12]

This approach divides metadata content for BDAs into three simple top-level groupings: Data, Reports, and Explorations. [13]

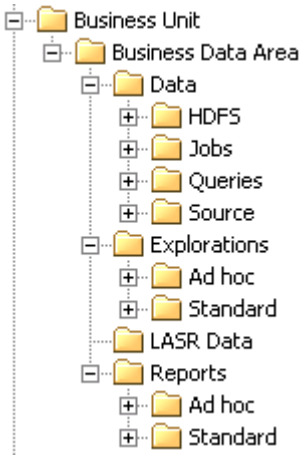

Figure 9. Example SAS Metadata Folder Structure

Business Unit

This is a container for the BDA folders. A user sees only those Business Data Area folders to which he or she has been given access. In addition, this can contain non-restricted data or content available to all members of the BU.

Business Data Area

This is an example of a BDA root-level metadata folder. This is the storage location for relevant metadata related to the BDA.

A BDA folder breaks down into the following:

- **Data—**This folder is visible only to data administrators for the BDA. This is the primary storage area for data-related metadata that is not LASR table metadata and therefore can be hidden from users. HDFS includes the library and table metadata related to HDFS. Jobs includes any jobs and scheduled jobs related to the movement of data into either HDFS or LASR. Queries contains saved queries built from the SAS Visual Data Builder. Source includes library and table metadata related to data source systems of the BDA.
- **Explorations—**The storage area for exploration-related metadata, only users with access to the SAS Visual Analytics Explorer see this folder. The Ad hoc subfolder is an area where power users are able to save their shared versions of their own content. The Standard subfolder contains exploration metadata generated through a standard development process. The business data administrator is able to move or create content in the Standard subfolder once it has gone through a release process.
- **LASR Data—**The storage area for all LASR libraries associated with the BDA SAS LASR Analytic Server and tables.
- **Reports—**The storage area for report-related metadata such as reports and images. The Ad hoc subfolder is an area where power users are able to save their shared versions of their own content. The Standard subfolder contains exploration metadata generated through a standard development process. The business data administrator is able to move or create content in the Standard subfolder once it has gone through a release process.

Further customer-specific subfolders are not shown in the approach example. For example, under **Reports→Standard** in one of the BDAs, you might segregate metadata in folders that organize the content by topic or project. The same applies to the Ad hoc area, where users have the freedom to build their own child folder structure to organize content.

TENANT DATA SEGREGATION

Data Security

You will likely have a requirement for strong security controls between tenants. As such, any data security restrictions between the tenants have to be preserved across potentially several layers of data.

Operating System and DBMS

SAS preserves the data security provided by your DBMS and operating system. It is essential to manage this physical layer access. For example, use host operating system protections to limit access to any sensitive data on the operating system. Use DBMS security controls to manage access to source data stored in the DBMS. It is important that any shared data access or load accounts across the user base are limited in scope and are not

shared between the tenants and BDAs (i.e., any service accounts are specific to a single BDA to prevent cross-boundary data access).

HDFS

If Hadoop is used as the co-located data provider, data stored in the HDFS needs to be segregated across tenants.

HDFS implements a permissions model for files and directories that are shared with the system used in UNIX. Each file and directory is associated with an owner and a group. The file or directory has separate permissions for the user that is the owner, for other users that are members of the group, and for all other users. For files, the Read permission is required to read the file and the Write permission is required to write or append to the file. For directories, the Read permission is required to list the contents of the directory, the Write permission is required to create or delete files or directories, and the Execute permission is required to access a child of the directory. Collectively, the permissions of a file or directory are its mode.

Each client process that accesses HDFS has a two-part identity composed of the user name and groups list. Whenever HDFS does a permissions check on a file or directory accessed by a client process, the following happens:

- If the user name matches the owner, the owner permissions are tested.
- Else, if the group matches any member of the groups list, group permissions are tested.
- Otherwise, other permissions are tested.

If a permissions check fails, the client operation fails.

SAS LASR Analytic Server

The SAS LASR Authorization Service resides in the middle tier. When a user performs an action within SAS Visual Analytics that requires access to a LASR table in the SAS LASR Analytic Server, the SAS LASR Authorization Service provides a signed grant to the application. This signed grant is generated based on the user's SAS metadata permissions on the resources involved. SAS Visual Analytics submits the signed grant to the SAS LASR Analytic Server. When the SAS LASR Analytic Server receives the signed grant, it performs the following actions:

- Validates the signature on the grant.
- Processes the request based on the constraints of the grant. For example, this might include limiting a user's access to rows via the permission conditions that apply to the data source.

The SAS LASR Authorization Service manages both table-level access and row-level access to data loaded into the SAS LASR Analytic Server.

In addition, for each SAS LASR Analytic Server, server signature files and table signature files are created on the operating system file system. These files need to be secured with OS controls, especially when data loading and unloading is allowed to happen outside of the SAS Visual Analytics web clients. It is possible for data administrators to circumvent the SAS LASR Authorization Service and SAS metadata permissions through code if these files are not secured.

The approach to secure these signature files is exactly the same as with any operation system files. It involves creating specific directories within a system path and then securing each directory appropriately using OS security controls. When a LASR library is defined in metadata or a SAS LASR Analytic Server is started through code, the signature location on the server can be configured. All activity for that particular server instance uses the specified subdirectory.

LOADING DATA INTO THE SAS LASR ANALYTIC SERVER

A simple way to collectively handle the security nuisances at each step in the data flow process is to follow these rules:

- Each BDA has its own SAS LASR Analytic Server defined that runs across the cluster.
- Each BDA has a unique service account identity that manages activities for the BDA's SAS LASR Analytic Server. (This allows OS and HDFS controls to be applied against the service account identity.)

In SAS Visual Analytics, passwordless SSH is used for internal communication and process execution between the machines in the SAS LASR Analytic Server cluster. Therefore, a working passwordless SSH configuration is required for the user account that completes the following actions:

- Start or stop a SAS LASR Analytic Server.
- Load, unload, append, or update data in a table in the SAS LASR Analytic Server.

Using Token Workspace to Link Service Accounts to SAS LASR Analytic Servers

SAS Workspace Servers interact with SAS by creating a server process for each client connection. Each SAS Workspace Server process enables client programs to access SAS libraries, perform tasks by using the SAS language, and retrieve the results. In the SAS Visual Analytics environment, tasks executed through the SAS Workspace Server are responsible for interacting with HDFS and the SAS LASR Analytic Server itself.

To facilitate these rules, specific BDA application server contexts need to be configured. These contexts contain a single SAS Workspace Server configured to use token authentication. The figure below shows what the generic setup would look like when turned into SASAppBDA application server context. In the actual implementation for a specific BDA, the BDA portion would be renamed to a BDA-specific identifier.

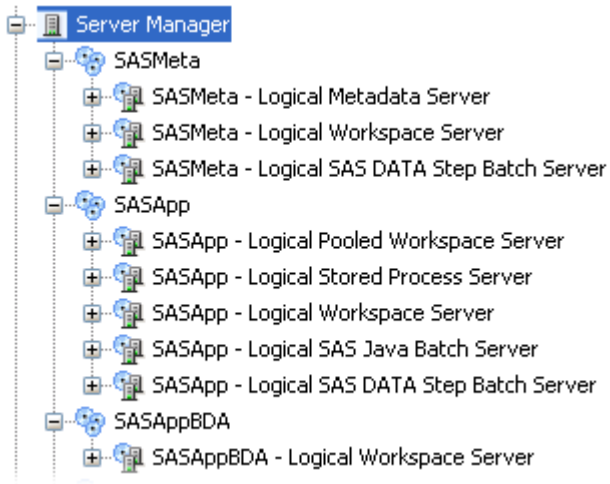

Figure 10. Application Server Context Configuration with a BDA-Specific SAS Workspace Server

Token authentication is when the SAS Metadata Server generates and validates a single-use identity token for each authentication event. This has the effect of causing participating SAS servers to accept users who are connected to the SAS Metadata Server. When used against a SAS Workspace Server, this means that no individual external accounts are required, no user passwords are stored in the metadata, and no reusable credentials are transmitted. So, as long as a user can log in to the SAS Metadata Server and has the required metadata access to resources, he or she can use token authentication to execute SAS processes and tasks using the BDA service accounts. As a result, service accounts associated with individual BUs or departments can be used behind the scenes.

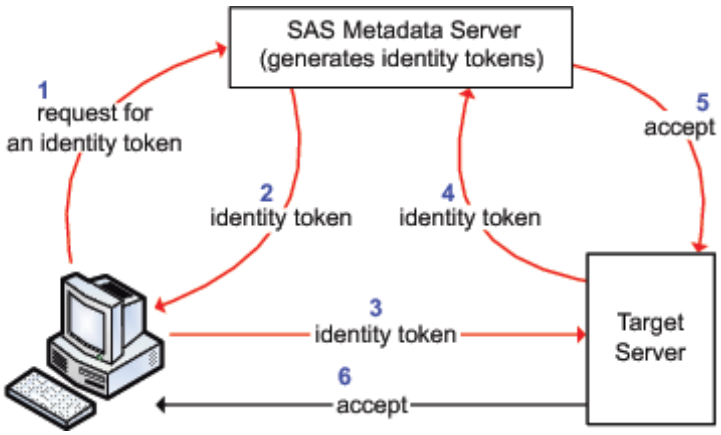

Figure 11. SAS Token Authentication

The BDA service accounts are used as the launch credentials of the token identity for each of their corresponding SAS Workspace Servers. These are made available to data administrator users (business and IT) to facilitate the interactive loading and unloading of data from source data sets into HDFS, starting and stopping in-memory LASR processes, and loading, unloading, appending, and updating data in the SAS LASR Analytic Server. For batch interactions in data in this scenario, the responsibility rests with central IT, who has access to all of the underlying service accounts and schedules through the normally invoked batch server belonging to the group-wide SASApp context. These service accounts are also the process owners of each BDA SAS LASR Analytic Server instance.

After this configuration has been put in place, users with the required metadata privileges and roles to administer data are allowed to select the relevant application server context for their tasks interacting with data in HDFS and the SAS LASR Analytic Server to execute in. This allows the security design to have the individual service accounts tightly integrated with OS permissions and Hadoop permissions to prevent cross-BDA data access. Access to launch processes belonging to the BDA service accounts is controlled through SAS metadata security and roles. Therefore, it is possible to have both business data administrators who can interact with their own BDA data and processes as well as IT-based data administrators who can interact with all BDAs.

The SAS Workspace Server, under the default SASApp application server context, has default host authentication, meaning that users cannot use this application server context to bypass security put in place due to the lack of any individual host accounts to launch processes. Even though SASApp can't be used for data loading, this application server context is left visible to facilitate SAS Visual Analytics standard functionality (i.e., retrieving map tiles for geo maps that use nonstandard SAS Workspace Servers like the SAS Pooled Workspace Server, which is just a platform-wide shared resource running under the out-of-the-box SAS General Servers account).

Lightweight Authentication Model

Thanks to the shared nature of the BDA service account that manages the flow of data into the SAS LASR Analytic Server in this distributed environment, a lightweight authentication model can keep the overhead of managing accounts between tenants low.

The basic premise of the authentication mechanism in this scenario is:

- Active Directory/LDAP Server is the authenticator for all standard users (no host accounts are required on the SAS LASR Analytic Server cluster).
- Host accounts that do exist are BDA service accounts.
- During the onboarding process for a new tenant, the BDA service account's passwordless SSH is configured and invoked behind the scenes through the use of token-based SAS Workspace Servers associated with those accounts.
- Access to invoke these token-based SAS Workspace Servers is limited to the central IT administration team to manage batch loading of data and to selected individuals within a BDA who have been designated as data administrators and allowed to load data interactively.

In SAS Visual Analytics distributed mode, all of the servers in the cluster are using an operating system based on Linux, so all servers share an authentication provider by default (the Linux host). The SAS LASR Analytic Server and the SAS Workspace Server always authenticate against the host, using whatever provider the host is configured to use. The SAS Metadata Server does the same by default, but can be configured to use an alternate mechanism.

To better handle the smooth addition and synchronization of business users to the SAS Visual Analytics environment, some of the authentication overhead, in terms of users and groups, are managed in an external authenticator. SAS Visual Analytics can be integrated with the company's Active Directory/LDAP Server to manage this. The integration level for SAS Visual Analytics will be the most lightweight through configuring the SAS Metadata Server to use direct LDAP authentication. In this case, the SAS Metadata Server validates users against an LDAP provider such as Active Directory. Direct LDAP enables the SAS Metadata Server to recognize accounts that aren't known to its host. Direct LDAP doesn't modify the host's behavior.

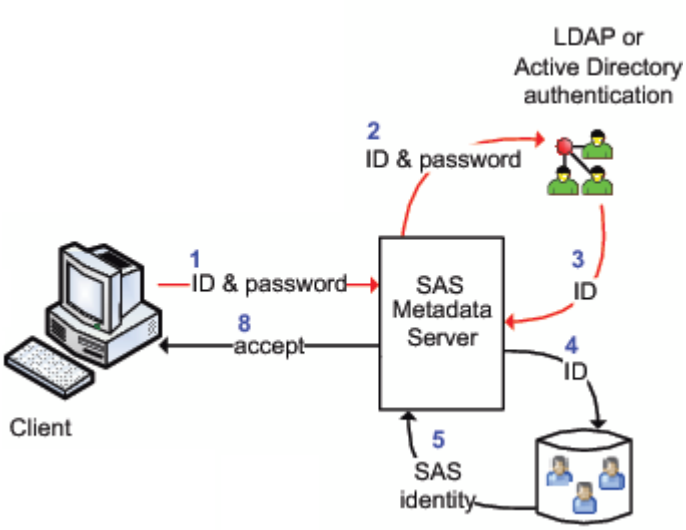

Figure 12. Direct LDAP Authentication

The benefits of this approach are:

- Using existing user identities in Active Directory/LDAP Server to authenticate users against the environment. With direct LDAP, you do not have to introduce or synchronize these identities across the SAS Visual Analytics cluster.
- Limiting the amount of host accounts and passwordless SSH configurations that need to be set up. Decreasing the maintenance required for standard users.
- Using existing IT processes for adding users to Active Directory/LDAP Server groups and potentially keeping SAS Metadata Server users synchronized with the central authenticator mechanism.
- The actual setup of this integration needs to happen only on the SAS Metadata Servers of the SAS Visual Analytics cluster.

SCRIPTING ONBOARDING TENANTS

All of these various steps to accommodate and segregate tenants are scriptable. This is an important aspect of minimizing the overhead and length of time to onboard a new tenant.

Operating system scripts are used to:

- Create the BDA service account on the host OS.
- Propagate account and SSH keys across the cluster.
- Create OS area for LASR signature files and secure it to BDA service account.
- Create the HDFS path for the BDA and secure it to the BDA service account.
- Launch the SAS Deployment Wizard to create a BDA SAS Application Server and BDA SAS Workspace Server.

SAS Open Metadata Interface and SAS Platform Object Framework are used to:

- Convert the BDA SAS Workspace Server to token authentication and associate it with the BDA service account.
- Define a BDA SAS LASR Analytic Server and library.
- Create the BDA folder structure based on the folder template.
- Create ACTs to secure the BDA folder structure.
- Apply the ACTs to the BDA folder structure and other resources.

SAS scripts are also used to manage the addition and synchronization of users and groups from the Active Directory/LDAP Server.

SCALABILITY

Scale Up

To scale up hardware is to add hardware resources such as CPU and memory to a node in the environment. In a VMware vCloud environment, this is relatively easy to achieve and can be done on all of the major components of a SAS Visual Analytics environment (SAS Metadata Server, SAS Compute Server, SAS middle tier, and LASR nodes). This is a little bit trickier in a commodity blade-based infrastructure because of the potentially limited expansion options available. It is worth noting that if scaling up is done on the LASR worker nodes, you should be aware that, assuming uniform data distribution and processing, the cluster performs at the speed of the slowest node. Any scaling up has to happen across all nodes or you will have unbalanced processing.

Scale Out

A more ideal option, especially for blade infrastructure, is scaling out. This is bringing in more nodes to increase workload. With SAS 9.4, all major components of the SAS Visual Analytics environment can be scaled out (with clustered metadata, SAS middle-tier, and load-balanced SAS Compute Server nodes). The same applies to SAS LASR Analytic Server nodes [14]. Adding additional hardware increases the overall data and computational capacity of the cluster. For the SAS LASR Analytic Server, this is a fairly linear (i.e., you double the amount of LASR worker nodes, you double the data and computational capacity).

CHARGEBACK AND COST MODEL

An important aspect of the provisioning of the private cloud is how to measure and chargeback usage to the tenants. There are generally two main options:

- Subscription—a recurring fee over a length of time.
- Usage based—track what is used with specific metrics that drive the bill.

For the customer scenario, a subscription model was chosen. The reason for avoiding a usage-based system was the likelihood that with pure usage-based costing (for example, direct CPU usage), there would be a determent of a wider adoption of the platform (i.e., if using something less directly means less cost, then this becomes the main driver of defining an everyday value for the adoption of the technology).

What can be used to determine the subscription cost in a SAS Visual Analytics environment for a tenant? Primarily, you must determine the capacity of the cluster and the selling portions of it to the tenants. These portions can be broken down as:

- Physical data storage—The platform has a finite amount of physical disk for the HDFS.
- In-memory data storage—The platform has a finite amount of memory to hold memory in the SAS LASR Analytic Server.
- Users—Based on the size of the cluster and the volume of the data, only a certain number can be supported at a time (for example, named and concurrent users). (This can be seen a quasi way of selling CPU capacity without actually tracking it.)

On the one hand, you have all of the platform costs—hardware, software, services, and maintenance. On the other hand, you have platform capacity. A base value can then be calculated and broken down into charge per GB for physical data storage, charge per GB for data in memory, and charge per user. The tenant chooses how much GB he or she would like to reserve on the physical disk, memory, and the number of users to produce a relevant subscription cost.

The model does not have to be totally cost neutral, it just needs to cover costs. A small overhead can be added, meaning that money could be accrued into an investment fund to start the process of scaling out the platform before capacity is completely full. This prevents problems such as when once capacity is reached, the next tenant has to effectively subsidize the scaling out and is expected to wait for the process to complete before onboarding.

CONCLUSION

SAS 9.4 and SAS Visual Analytics support many cloud deployment options to match your specific needs. The paper demonstrates the processes and design that go into building a truly enterprise cloud platform that can reveal the value locked in data to all areas of your business.

REFERENCES

[1] Mell, Peter, and Grance,Timothy. 2011. “The NIST Definition of Cloud Computing”. *National Institute of Standards and Technology*. Available at: http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf.

[2] “SAS® 9.4 Support for Web Browsers and Plug-Ins.” Available at:
http://support.sas.com/resources/thirdpartysupport/v94/browsers.html.

[3] “SAS® Visual Analytics for SAS® Cloud: Quick-Start Guide.” Available at:
http://support.sas.com/documentation/solutions/cloudgen/9.4/VACLD_Quick_Start_Guide.pdf.

[4] “Tagging Your Amazon EC2 Resources.” Available at:
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Using_Tags.html.

[5] Deployment and administration documentation. Available at:
http://support.sas.com/documentation/onlinedoc/va/index.html.

[6] VMware vCloud Architecture Toolkit (vCAT) 3.1.2 at:

http://www.vmware.com/cloud-computing/cloud-architecture/vcat-toolkit3.html

[7] “Enabling a CPUfreq Governor.” Available at:

https://access.redhat.com/site/documentation/en-US/Red_Hat_Enterprise_Linux/6/html/Power_Management_Guide/cpufreq_setup.html#enabling_a_cpufreq_governor.

[8] “Make the Most of Your Analytical Talent.” Available at:
http://www.sas.com/offices/europe/uk/resources/whitepaper/wp_20417.pdf.

[9] John R. Rymer’s Blog. “Understanding Cloud's Multitenancy.” Available at:
http://blogs.forrester.com/john_r_rymer/12-03-20-understanding_clouds_multitenancy.

[10] Macfarlane, Andrew, and Schneider, Frank. 2013. “Pi in the Sky: Building a Private SAS® Cloud.” *Proceedings of the SAS Global Forum 2013 Conference*. Cary, NC: SAS Institute Inc. Available at
http://support.sas.com/resources/papers/proceedings13/.

[11] Gahagan, Ken. 2014. “SAS® Grid Manager, SAS® Visual Analytics, and SAS® High-Performance Analytics: Sharing Hardware and More.” *Proceedings of the SAS Global Forum 2014 Conference*. Cary, NC: SAS Institute Inc.

[12] SAS® 9.4 Intelligence Platform: Security Administration Guide, Second Edition. Available at:
http://support.sas.com/documentation/cdl/en/bisecag/67045/HTML/default/viewer.htm.

[13] Chitale, Anand, and Redpath, Christopher. 2013. “Whirlwind Tour Around SAS® Visual Analytics.” *Proceedings of the SAS Global Forum 2013 Conference*. Cary, NC: SAS Institute Inc. Available at
http://support.sas.com/resources/papers/proceedings13/

[14] SAS LASR Analytic Server 2.2: Reference Guide. Available at:
http://support.sas.com/documentation/onlinedoc/securedoc/index_lasrserver.html.

ACKNOWLEDGMENTS

We would like to acknowledge the great work and contributions to this paper by Rob Stephens and Erwan Granger.

CONTACT INFORMATION

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

Christopher Redpath
SAS Institute Inc.
100 SAS Campus Drive
Cary, NC 27513
christopher.redpath@sas.com

Nicholas Eayrs
SAS Institute Inc.
100 SAS Campus Drive
Cary, NC 27513
nicholas.eayrs@sas.com