

SAS® Viya®: Experience Gained through Multiple Installations and Implementations

John Schmitz, Luminare Data LLC

ABSTRACT

SAS® Viya® offers new and exciting capabilities for the SAS® Platform, but the SAS Viya capabilities and designs might be unfamiliar to many SAS® users. This paper highlights experiences gained from multiple SAS Viya installation and implementation projects involving SAS Viya 3.2, 3.3, and 3.4, from small to very large hardware resource allocations. Key topics include hardware and architectural concepts, leveraging of external data sources, recovery considerations, integration with SAS® 9, and daily operation overviews. This paper provides valuable insights to organizations considering or in the early stages of a SAS Viya implementation.

INTRODUCTION

SAS® Viya® is the latest modernization of the SAS platform. It offers new and exciting capabilities, especially around 'Big Data' and machine learning. Industry interest in this product appears to be high, but few SAS users have direct experience with the software and its features. The goal of this paper is to share some of the direct experiences gleaned from multiple Viya installations and implementation efforts.

SAS Viya is primarily built to support memory resident, advanced analytics and machine learning across clustered machines. This is a significant architectural departure from the SAS 9 platform and is better matched to the current 'Big Data' approach to modeling. That being said, it is not as far removed from the traditional SAS environment as one may initially think. This is an important consideration since the experienced SAS user has a head start in understanding this platform. This is not necessarily the case for the SAS administrator who will encounter a significantly different architecture for implementation and management.

This discussion attempts to highlight similarities, differences and unique strengths of the two platforms. The information shared is derived from the authors experience and opinions generated from installations, implementations and usage of SAS Viya that includes version 3.2, 3.3, and 3.4, as well as first hand conversations with others who have been involved with similar efforts. Some experience has been generated via laboratory testing to evaluate certain capabilities while others involve business-grade production deployments encompassing multiple terabytes of RAM and hundreds of CPUs, attached to large scale Hadoop clusters with similar resource allocations.

This paper assumes the reader is generally familiar with the key features and architecture of SAS 9. Since SAS Viya is predominantly deploy in Linux systems within the RHEL family, all comments are written from the RHEL operating system perspective and will discuss SAS Viya 3.4 unless otherwise stated. This paper focuses on core elements of interest to most who would leverage the platform. There are many additional elements and capabilities that can leveraged within the SAS Viya implementation that are not addressed here.

WHAT IS SAS VIYA

SAS Viya should not be considered a replacement for SAS 9, but rather an integrated, complementary, modernization to the SAS Platform. Viya works together with SAS 9 to expand the analytics capabilities of the SAS Platform. In order to provide this

modernization, Viya differentiates itself from SAS 9 in several areas. For this discussion, the differentiators are categorized into three distinct areas:

- The CAS Engine
- The Microservice Components
- The SAS Viya Hardware Architecture

The discussion here is not meant to provide an exhaustive list of differences, but rather to highlight and add clarity to some of the core features where SAS Viya generates significant enhancements and changes compared to the SAS 9 environment.

THE CAS ENGINE

The CAS engine is arguably the most significant difference between SAS 9 and SAS Viya. The CAS (Cloud Analytic Service) engine provides a data and analytic framework that can be executed across any number of distributed machine nodes, limited by the number of CPUs licensed for the Viya installation.

CAS both integrates and enhances the capabilities of LASR and HPA in SAS 9. LASR is the distributed, in-memory data engine that supports Visual Analytics. HPA (high performance analytics) is the engine built to support distributed, parallelized analytics processes used to reduce runtimes for advanced analytics procedures. CAS integrates their functionality so the same engine and data can be used for both the analytics and machine learning algorithms, as well as the reporting and analytic components provided through Visual Analytics. Importantly, CAS is built to handle the analytics workload required to support large-scale machine learning and artificial intelligence.

CAS enhances these engines through increased fault tolerance and extensions beyond in-memory data. CAS is designed to withstand node failure during the analytics run without requiring a job restart. Data are distributed across nodes to ensure data redundancy by requiring data appear on two or more machines in the cluster. The CAS engine also allows for data swapping in the event that memory requirements exceed physical RAM. This swapping feature is different than SASWORK. In CAS, swap files are memory-mapped data accessible to the CAS engine as needed. These memory mapped files can be swapped in or out of RAM so that the engine can support data beyond the physical RAM limits. CAS also supports a file storage format (SASHDAT) which can be used to efficiently store and retrieve in-memory CAS data tables.

THE MICROSERVICE COMPONENTS

The SAS 9 enterprise architecture is monolithic design that is highly integrated. This design increases the complexity of system upgrades, enhancements and scaling since the entire application suite must be considered when implementing a system change.

SAS Viya replaces the monolithic design with a microservice architecture. Primarily, these services replace the SAS 9 midtier and metadata server components. The microservice 'loosely coupled' design allows greater flexibility since most services can be managed with a greater level of independence. Individual services can be restarted, updated, or relocated for better scalability, without requiring changes to other services. Within the microservice architecture, each service can be designed with a specific functionality in mind. This can greatly reduce the effort to correct issues or expand capabilities. This new architectural approach should simplify the update and maintenance of the Viya platform, compared to the SAS 9 installation. It also allows greater flexibility within the organization to ensure a better match between system load and assigned resources. This design does have an associated cost in terms of an increased need for logging and communication to support services communicate across the Viya platform.

THE SAS VIYA HARDWARE ARCHITECTURE

Without delving too deeply into the machine architecture, it may be helpful to highlight the core components that make up this platform. The SAS Viya installation can exist on a single machine but is commonly installed across a multiple-machine cluster. This clustered implementation is not like SAS Grid where separate SAS 9 installations execute jobs on an assigned machine or execute independent steps that can execute simultaneously across multiple machines. Rather through the CAS engine, the individual step can execute across multiple machines in a true parallel process with each node responsible for the portion of data residing on that machine.

An environment diagram from *SAS VIYA 3.4 Administration* (SAS Institute Inc, 2018) is reproduced for reference in Figure 1. The typical components within this architecture include:

- CAS Services
 - o CAS Controller
 - o CAS Secondary Controller (optional)
 - o CAS Workers
- SAS Services
- SAS Programming
- Ansible / VIRK Support
- LDAP Server Support

The first three areas provide the core functionality for SAS Viya while the latter two are not shipped as part of the software deployment but are used to support the Viya installation.

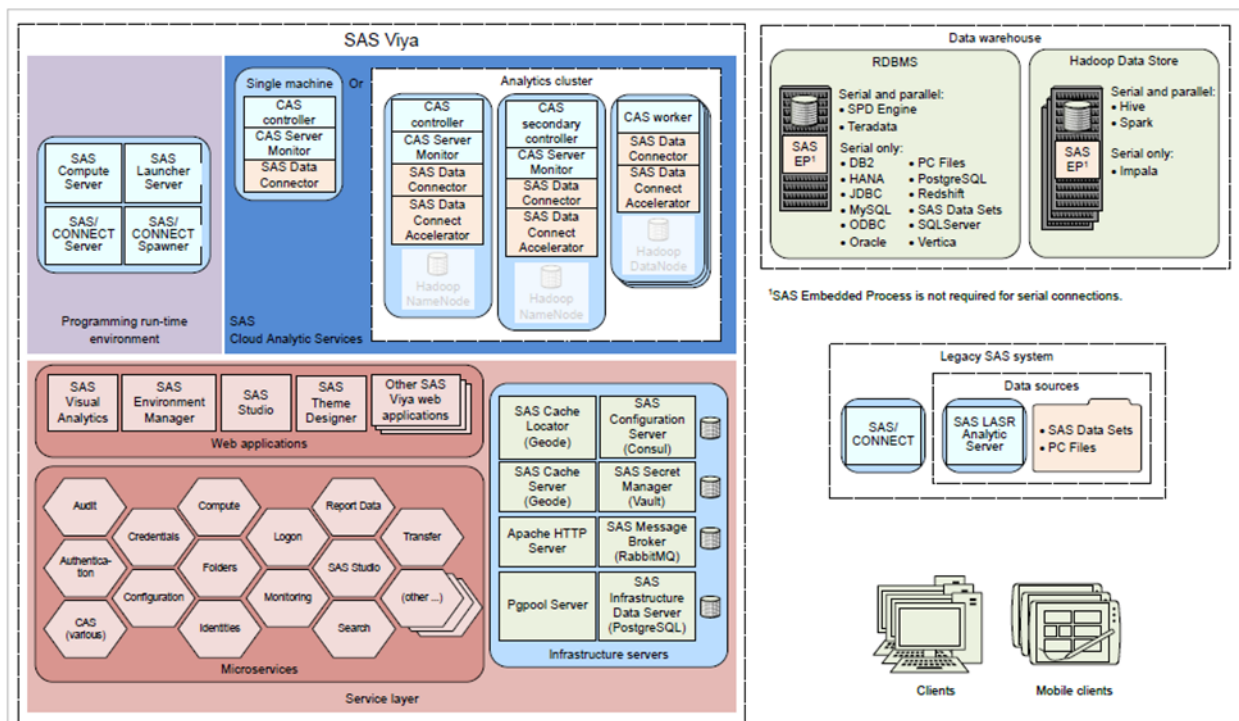


Figure 1. SAS Viya 3.4 Full Deployment Diagram.

CAS SERVICES

The CAS Engine operates on the CAS controller and any CAS worker nodes that may be present within the SAS Viya installation. The CAS controller can simultaneously serve as a worker node, but the roles of CAS controller and CAS worker are distinct. In Figure 1, the CAS service facilities are represented in the upper left section. While the CAS engine is responsible for handling the analytic workloads, these CAS service machines execute the CAS requests, making them the underlying workhorse of the SAS Viya deployment. CAS distributes data across the worker machines, allowing each machine to process requests based on the data distributed to that particular host. Data redundancy across hosts is used to provide fault tolerance should any worker node fail while processing the CAS request.

The CAS controller is critical to CAS operations and could represent a single failure point. Starting with SAS Viya 3.3, the CAS Services can include an optional secondary CAS controller. This machine is essentially a redundant hot backup that can assume the role of CAS controller in the event of a primary controller failure. However, the two controllers do not share load or jointly operate as controller. Should a failure on the primary controller occur resulting in a switch to the secondary controller, administrator intervention is required to return the operations to the primary controller.

SAS SERVICES

Microservices operate on the SAS Services node. In Figure 1, the various microservices are represented by the lower left section of the graphic. These services can be installed on the same machine as CAS controller or be installed on one or more machines dedicated to support microservice operations. It is possible to provide resource scaling of these services by spreading them across multiple machines or establish horizontal scaling by installing the same service to multiple machines.

Various infrastructure services are also required in a SAS Viya implementation and would typically share hardware with microservices. These infrastructure servers include message brokers, data servers and more that facilitate the logging and messaging across the various microservices.

A basic testing environment will need at least 64gb RAM to install the various microservices for a reasonably full suite of products, assuming the machine supports SAS Services only. When working with more limited hardware resources, it is also possible to distribute these services across four machines with 16gb RAM each. Additional RAM beyond the 64gb baseline will be required to support more users and larger work efforts.

SAS PROGRAMMING

The programming node is conceptually separate from each of the above; although it will commonly share hardware with the CAS controller node. In Figure 1, the programming services are shown in the upper left section. The programming node supports SAS Studio and other SAS V9 environment elements including SAS Foundation, SAS compute server, SAS workspace server, SAS connect server and support for licensed SAS/ACCESS® engines.

The SAS programming features require different user authentication methods than the remaining Viya elements. SAS Viya services primarily run under service accounts, allowing the system to manage access to resources without requiring users to have login accounts on the underlying hardware operating systems. The SAS Programming environment operates under a different authentication method and will typically require login accounts on the host system for users that directly leverage this feature set.

ANSIBLE / VIRK SUPPORT

Neither Ansible or VIRK are part of the SAS software deployment for SAS Viya or represented in Figure 1, but both are used to support the multiple system architecture.

Ansible is an open-source tool primarily involved in the installation, configuration and on-going maintenance of SAS Viya. The Ansible installation can reside on any of the Viya cluster machines or on another system that has access to these systems.

As part of the software licensing, SAS provides an Ansible playbook which is user-customized to meet the desired installation architecture. Once properly customized, Ansible executes the playbook to deploy the software.

SAS Viya does not require a SAS Software Depot like SAS 9. Rather the software is deployed using rpm packages and yum. However, a local repository can be created to ensure all machines and updates remain in-sync during updates.

VIRK (Viya Infrastructure Resource Kit) is a set of tools made available by SAS through GITHUB at <https://github.com/sassoftware/virk>. VIRK is essentially a collection of Ansible playbooks so the files should be placed wherever the Ansible installation resides. The VIRK tools include pre-deployment assessment scripts as well as tools that perform various administrative functions.

One of the handy tools offered through VIRK is a series of scripts to start, stop, and status check all services within the Viya installation. It provides a single place to trigger commands that will access the services across all machines in the deployment. Since Ansible has access to each machine in the installation, VIRK can properly sequence START and STOP activities across the cluster and simplify administration across larger installations.

LDAP SERVER SUPPORT

The LDAP (Lightweight Directory Access Protocol) service replaces functionality from the SAS metadata server in the SAS 9 environment. LDAP is commonly used in Linux environments to manage user login credentials throughout the network. SAS Viya leverages LDAP directly to manage user credentials.

SAS VIYA INSTALLATION

The SAS Viya installation experience is notably different than that used by SAS 9. The SAS Download Manager, SAS Software Depot, SAS Deployment Wizard, and installation plans are not used in Viya. Rather, the install primarily uses the Ansible installation playbook included with the software licenses. A local repository of the rpm files can be established which may help systems keep services in-sync, but it is not required.

Although Viya can be installed manually using YUM, Ansible is the recommended installation practice. When Ansible is used, the installation location for the various components are determined by the inventory.ini file which is included as part of the SAS-generated playbook. The standard file installs all components on a single machine and would not include a backup controller or separate worker nodes. User customizations include listing machines to use for Viya and assign the various components to specific machines as needed. The Ansible installation process also leverages SSH, preferably passwordless SSH, and sudo command access to complete the installation.

SAS VIYA AND HADOOP

Many companies that are pursuing Big Data analytics efforts maintain a Hadoop cluster. For such organizations, there are important opportunities and considerations for the SAS Viya installation and implementation efforts.

CO-LOCATION

The CAS installation can be co-located on the Hadoop data nodes so that the same hardware supports both Hadoop and Viya. This option can simplify connectivity and increase throughput between the two systems with a trade-off of increased resources contention on the supporting hosts. Within this deployment, the CAS controller node would reside on the Hadoop NameNode while the CAS workers would reside on Hadoop DataNodes. It is possible to install workers on a subset of the Hadoop data nodes as well as to allocate a reduced set of CPUs from the DataNodes to the CAS cluster.

The co-located deployment benefits from faster connectivity between workers from both systems. This can be very beneficial when SASHDAT files are stored within HDFS.

EMBEDDED PROCESSING

Organizations that choose not to co-locate Hadoop and Viya may still benefit from the Hadoop embedded process. The embedded process allows Viya to push analysis directly to the Hadoop installation for processing. The embedded process can also support parallelized read/write activity directly between Viya worker nodes and Hadoop work nodes, decreasing the time to transfer data. Parallelized loads can be used to quickly restore data in event of a system outage or the need to reload data required to support specific work requests.

The embedded process is also leveraged to support in-database processing for Hadoop. In this case, the CAS data requests are passed to Hadoop to process. Hadoop can complete many of the analytics request through native system processes, potentially eliminating the need to transfer data from Hadoop to CAS or load the data into memory. This can lead to simplified data management and reduced demand on in-memory resources.

SAS/ACCESS® FOR HADOOP

A third option for Viya/Hadoop connectivity is the SAS Data Connectors and SAS/ACCESS® for Hadoop engines, leveraging the SAS 9 access engine implemented through the SAS Programming environment. This implementation allows transfer of data between SAS Viya and the Hadoop system's HIVE data as a serial data transfer only. Within this process, a serial transfer is generated between the CAS controller node and the NameNode on the Hadoop system. Hadoop gathers all the required data from across the various data nodes on the cluster and transfers this data back to Viya. The CAS controller receives the data and distributes across the various CAS worker nodes on the cluster. The transfer relies heavily on the two systems' primary nodes and the network path between them, making the transfer dependent upon resources availability on these critical elements.

SAS VIYA AND OTHER DATA SOURCES

SAS 9 is highly regarded for its ability to interface with a wide variety of data sources, and SAS retains much of that capability within Viya. The upper right section of Figure 1 shows some of the data interfacing capabilities of the Viya platform. Support for different data sources can vary, depending on licensed products and the specific source. SA Viya supports several SAS-based sources including:

- SASHDAT
- LASR Analytics Server
- SPDE (SAS Scalable Performance Data Engine)
- SAS data sets

Other data sources may be supported with:

- In-database processing which would support parallel data transfers
- SAS data connectors which support serial transfers between the data source and CAS
- SAS 9 integrated access which leveraging the Programming Services to transfer data into the SAS Foundations environment.

SAS Viya can access data from most common database systems. The most current information on supported databases can be found in the *SAS® Viya® 3.4 for Linux: Deployment Guide* (SAS Institute Inc, 2018).

SAS VIYA WEB APPLICATIONS

Web Applications are depicted in Figure 1 in the lower left section and typically reside on the same hardware as the microservices. Unlike SAS 9, Viya does not ship with JAVA-based clients such as SAS Enterprise Guide®, SAS Enterprise Miner™, SAS Management Console, or SAS Data Integration Studio client. Rather, access is established through various web applications. Since SAS Viya offers multiple product licensing options, specific web applications will vary by installation. A sample listing of web applications from a more complete installation can include:

- SASBackupManager
- SASDataExplorer
- SASDataStudio
- SASDecisionManager
- SASDrive
- SASEnvironmentManager
- SASGraphBuilder
- SASJobExecution
- SASLineage
- SASLogon
- SASModelManager
- SASQKBManager
- SASReportViewer
- SASStudio
- SASStudioV
- SASThemeDesigner
- SASVisualAnalytics
- SASWorkflowManager

Some of the more commonly installed components include two versions of SAS Studio, Visual Analytics, Model Studio and Environment Manager.

SAS STUDIO AND SAS STUDIOV

SAS Viya includes SAS Studio (Figure 2) as one option to interface with the programming environment. SAS Studio is a browser-based application derived from the SAS Studio shipped with the SAS 9 environment and operates in a similar fashion. It leverages many of the SAS 9 elements included with the programming node including the SAS Workspace server but does not support the visual programming model of the SAS 9 version. SAS Studio requires users to have a login account on the host systems since the Workspace server processes run under the individual user login. This login requirement is unique to SAS Studio. Other processes run under system accounts so user login account at the operating system level would not be required for other components.

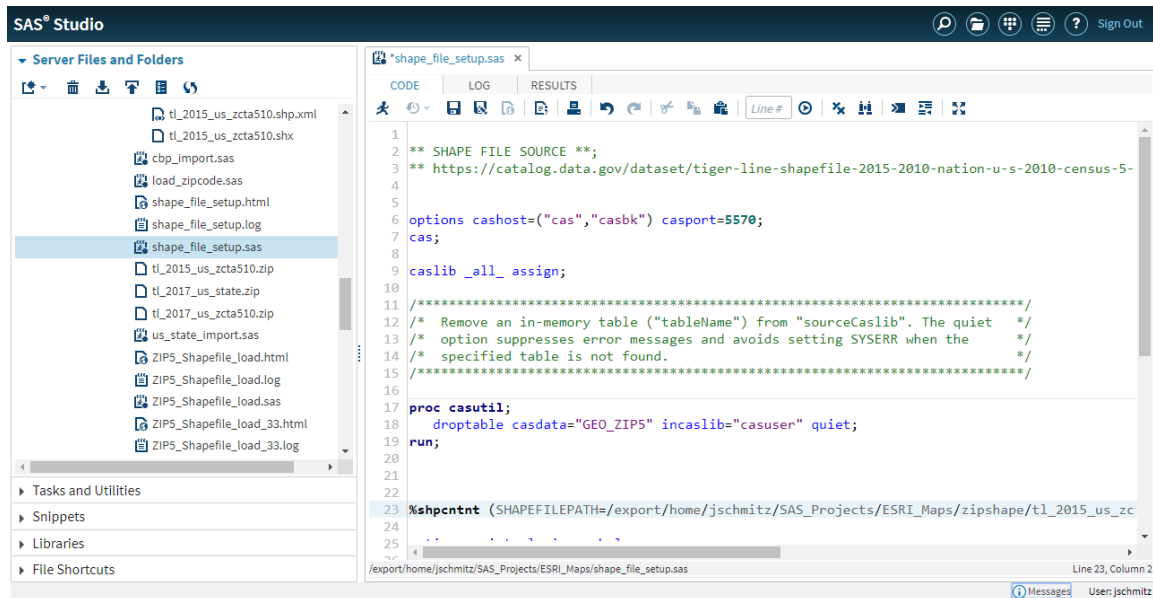


Figure 2. SAS Studio Web Environment Screenshot.

SAS Viya 3.4 introduced SAS StudioV (Figure 3) as an additional programming environment. This new environment is critically different than its predecessor. Unlike the predecessor environment, StudioV is built as a microservice. This allows the newer editor to operate under the same authentication mechanisms and generic service account as the other Viya components. This change eliminates the need for individual logons to the host system for those who use StudioV. However, one important consequence is that StudioV does NOT access the operating system home directory or other folders for opening and saving of SAS programs. Rather, it relies on SASDrive as a program storage facility. For environments where file system access is required or retained as a capability, SAS Studio support and hence the user logon capabilities will remain a requirement. Note that both environments can be implemented in the same environment without issue.

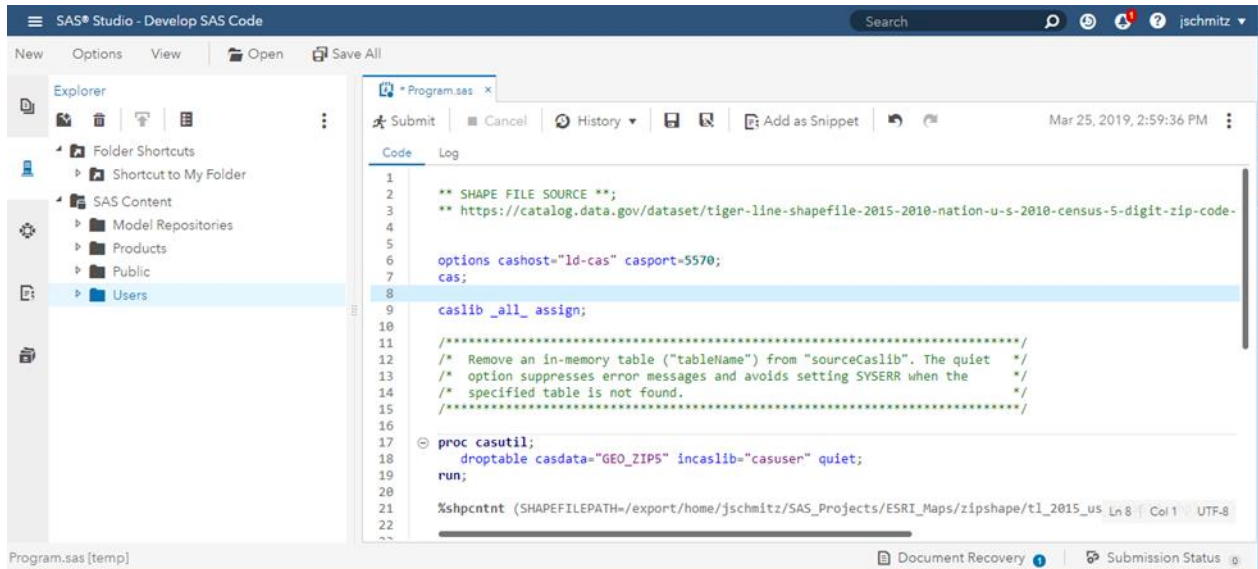


Figure 3 SAS StudioV Web Environment Snapshot.

SAS VISUAL ANALYTICS

SAS Viya offers a newer version of SAS Visual Analytics as a foundational user interface to access many of the environment’s capabilities within a visual GUI interface (see Figure 4). Visual Analytics is heavily integrated into the Viya experience, with most of the related visual tools accessible through it. Unlike SAS 9 where Visual Analytics requires loading a copy of the data into LASR, the Viya-based version leverages CAS, leading to a more synergistic integration with the remaining components.

The Visual Analytics application provide extensive reporting capabilities within the Viya platform, but it also provides a natural springboard into other components. The Objects

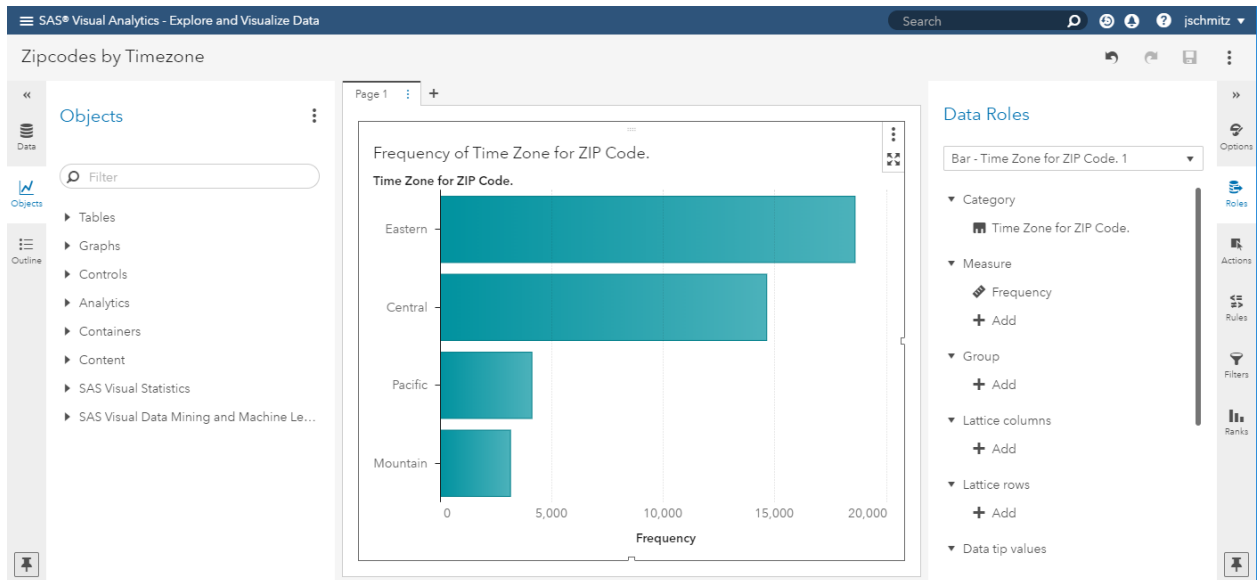


Figure 4 SAS Visual Analytics Web Application Screenshot from SAS Viya 3.4.

section also shown in Figure 4 shows other features available including Analytics, Visual Statistics, and Visual Data Mining and Machine Learning. Some of the standard and optional analytics components within these groups include:

- Analytics
 - Automated Analysis
 - Forecasting
 - Network Analysis
 - Path Analysis
 - Text Topics
- SAS Visual Statistics
 - Cluster
 - Decision Tree
 - Generalized Additive Models
 - Generalized Linear Model
 - Linear Regression
 - Logistic Regression
 - Model Comparison
 - Nonparametric Logistic Regression
- SAS Visual Data Mining and Machine Learning
 - Bayesian Network
 - Factorization Machine
 - Forest
 - Gradient Boosting
 - Neural Network
 - Support Vector Machine

The Data section allows access to various data sources and the ability to create derived measures that may be required by different reporting and analytic procedures.

SAS MODEL STUDIO

SAS Model Studio provides a wealth of modeling tools within Data Mining and Machine Learning models, Forecasting and Text Analytics, leveraging data sets loaded into CAS. The product includes common modeling tools to build and estimate and compare alternative models.

The details of these models are beyond the scope of this discussion, but the screenshot in Figure 5 provides a sample model pipeline example. Analysts can use Model Studio to create pipeline jobs that integrate data management, statistical and machine learning algorithms, and model comparisons. Its capabilities would be most closely related to Enterprise Miner within the SAS 9 platform.

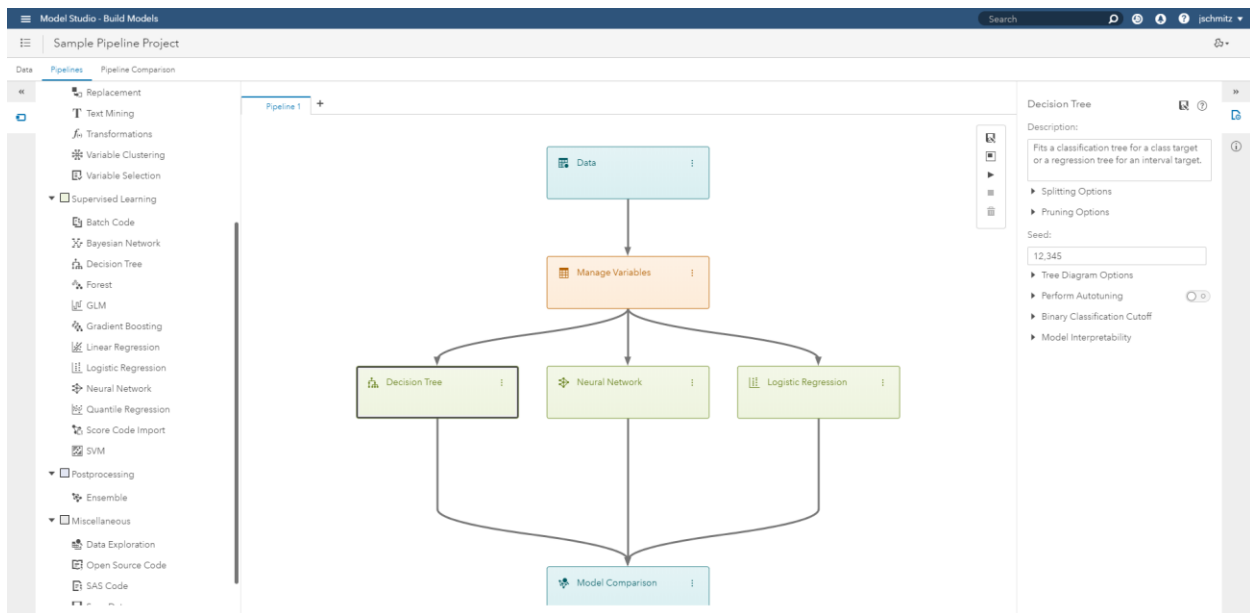


Figure 5. SAS Model Studio Pipeline Screenshot.

SAS ENVIRONMENT MANAGER

The SAS Environment Manager (Figure 6) is the Viya administrator’s primary location to manage the platform. It offers details on system resources, active sessions, microservice configurations and more. It is notably different than the SAS 9 Environment Manager and encompasses the capabilities of SAS 9 Environment Manager and Management Console. The Dashboard’s left-hand graphic displays a color-coded status for each machine and service in the installation. The right-hand graph shows memory utilization for each CAS machine or can display CPU load by machine.

The left-hand panel includes multiple shortcuts to the various capabilities of the Web Application. These shortcuts open pages to allow administrators access to the various components for the system. These shortcuts provide access to additional information on machine utilization, logs, running tasks, user management, system configuration and other items.

SOME LESSONS LEARNED

HADOOP FOR MEMORY-RESIDENT DATA BACKUP

Reloading memory-resident data into the CAS engine following a system outage can be a daunting effort. This can be especially true if the data are not stored since that may require rerunning the data queries that originally loaded the data to CAS.

SAS Viya can utilize HDFS to store data in a memory-ready format, while placing the data in the larger storage space allocated to the Hadoop cluster. These files can be used to quickly reload data into the in-memory CAS engine, either as recovery support or used to load data that may not be available on an on-going basis. When the embedded processing capabilities are licensed, Viya can leverage parallel data load from Hadoop to the CAS in-memory systems greatly reducing the time required to reload data tables.

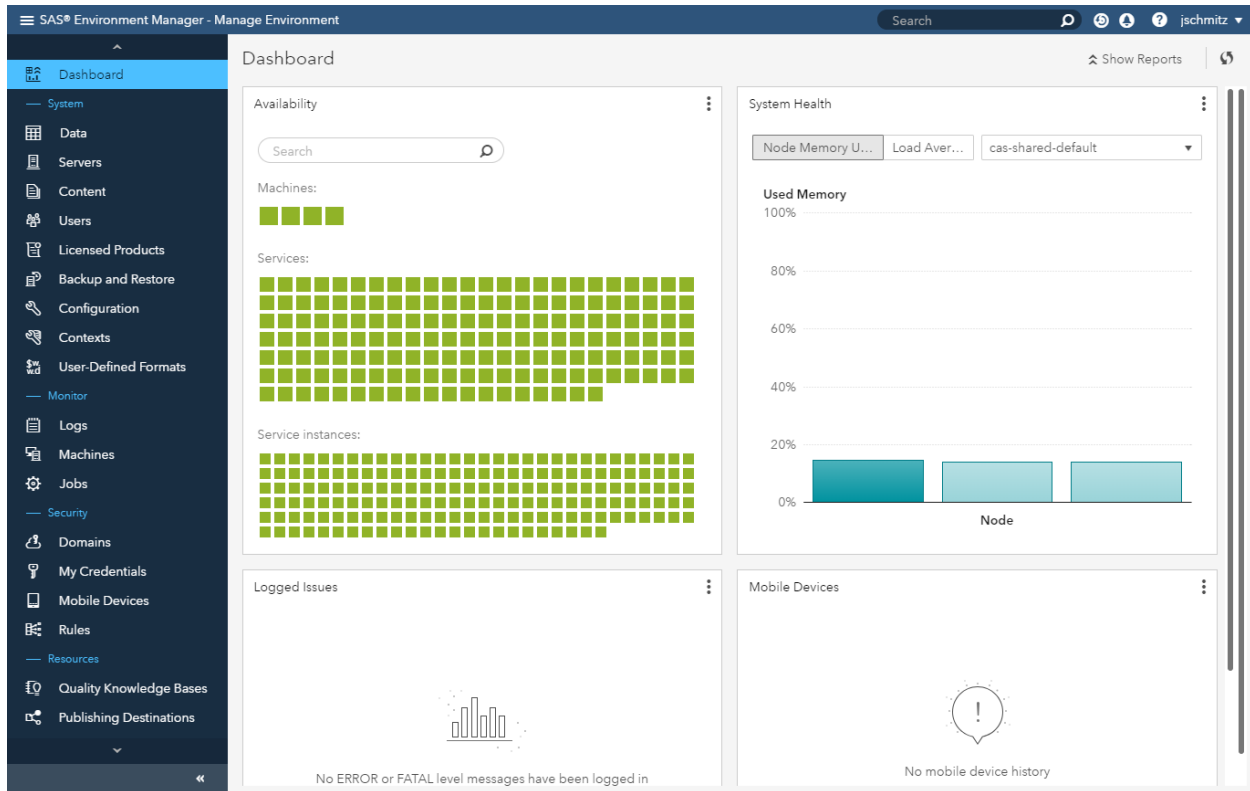


Figure 6. SAS Environment Manager Dashboard and Screenshot.

SYSTEM SHUTDOWN / RESTART

System shutdown and restarts can become corrupt when executed incorrectly, making future restarts more difficult. The overall process seems more reliable in SAS Viya 3.4 than in 3.3. It also appears that the VIRK tools can help alleviate these issues since it can coordinate start and stop activities across systems. The VIRK tools appear more critical as the number of unique machines in the installation increases.

ISOLATE RESOURCES ON LARGER INSTALLATIONS

Larger installs will spread across multiple machines. Experience with 25 node systems suggest that installations of this size will likely benefit from a system dedicated to the controller node role with no additional installed components. The combined load of the CAS controller, microservices, web applications, and infrastructure servers can overload even a large machine when the system becomes load with multiple users and projects. The optimal tweaking depends heavily on the underlying hardware resources allocated to the system and the specific deployment use cases. However, a good first cut appears to be separating the CAS controller from the microservices and infrastructure servers. A second cut may be to split microservices and infrastructure servers to separate hardware.

CAS BACKUP CONTROLLER CONSIDERATIONS

Laboratory evaluations of the secondary CAS node on SAS Viya 3.3 and 3.4 have been generally favorable when high availability is necessary. These experimentations have confirmed that that the backup controller can keep the system running in the event of a failure of the main controller. One should note there has been reports of issues with

upgrades from 3.3 to 3.4 when a backup controller was included in the 3.3 installation. It is strongly advised that user who intend to use a secondary controller move to Viya 3.4.

There is an important trade-off since the CPUs allocated to the backup controller count as part of the total CPU count in the cluster. Hence a backup controller may cost an additional worker node of similar size in the hardware mix.

BROWSER CACHE CAN LIMIT LARGE-SCALE VISUALIZATIONS

Even though SAS Viya can handle massive data series, the visualization tools can be limited by the cache capabilities of the browser. User experience with larger choropleth maps has shown that the browser cannot necessarily display the number of regions that were desired on a map. In other cases, geoplots and scatter plots were necessarily limited in scope to manage the capabilities of the browser to display and manipulate data.

The Viya administrator can modify defined query limits to expand the plotting capabilities beyond default levels. This requires a level of caution since browser instability and poor browser performance will result when these defined thresholds are increased too far. Reasonable adjustments that are tailored to the business needs can be made. However, more extreme cases will result in users needing to experiment with alternative visualization methods, refining granularity, and adjusting filters to adequately display the desired results.

LEVERAGING SAS 9

For organizations with SAS 9.4, it is possible to access Viya from their standard SAS 9 programming environment. Starting with SAS 9.4M5, SAS 9 can directly connect to the CAS engine in Viya and access the capabilities of the CAS engine. Properly leveraged, this feature adds extensive power and capabilities to the SAS 9 platform. CAS-specific commands can be added to the SAS 9 code and executed on the Viya platform, while data uploads and downloads can be managed simply. Although SAS Viya does not ship with Enterprise Guide, users can continue to use it as a development tool by connecting to a CAS engine and developing code as they would in SAS Studio. Older SAS version can still leverage the approach, but older version must connect to SAS Viya via SAS/CONNECT creating an alternative, but less seamless, approach to access the Viya platform.

ACCEPTANCE BY DATA SCIENCE TEAMS

Integration and acceptance by the users are always critical elements of any software implementation. The author has had first-hand experience with Data Science teams using the software in business application. The reaction can vary, especially depending on the user's past software experience.

As one would expect, those who have used SAS 9 heavily are among the quickest to embrace Viya. SAS Viya receives high marks for its speed and extensive capabilities, but these users also notice the absence tools such as Enterprise Guide, Forecast Studio, and Enterprise Minor.

R users who have the platform recognition its usability and power of SAS Viya. One such user has stated that the visual interface is intuitive, and the analytics power exceeded what he had seen with R, prompting him to invest the time to learn more with the idea that Viya would become his environment of choice.

CONCLUSIONS

SAS Viya is the latest modernization to the SAS platform, especially extending the platform's capabilities in big data. SAS users will find new capabilities in machine learning with the processing power to support large-scale analytic endeavors. These users will find

some familiarity with SAS 9 accompanied by new implementations and algorithms. The SAS administrator will experience a radically different installation and administration environment.

SAS Viya offers data science teams a robust set of analytic tools and can perform well under real-world business load. SAS Viya 3.2 provides notably fewer capabilities than either 3.3 or 3.4. Analysts were able to generate solid analysis leveraging SAS Viya 3.3. Viya 3.4 has expanded these capabilities, improved resource utilization and polished some features to create a cleaner, more resilient platform. The tighter integration between the CAS engine and SAS 9.4M5 can simplify established SAS users' transition between the two platforms. SAS Viya, especially with its 3.4 release has demonstrated acceptance by advanced users for its capabilities but lacks some of the more user-friendly tools found in SAS 9. With Viya 3.3 and more recently with 3.4, the SAS Viya platform appears ready to handle the demands of real-world analytics.

REFERENCES

SAS Institute Inc. 2018. *SAS® Viya® 3.4 Administration*. Cary, NC: SAS Institute Inc.

SAS Institute Inc. 2018. *SAS® Viya® 3.4 for Linux: Deployment Guide*. Cary, NC: SAS Institute Inc.

CONTACT INFORMATION

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

John Schmitz
Luminare Data LLC
john.schmitz@luminaredata.com
Luminaredata.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.