Paper 4947-2020

Evolve Your SAS® Administration Skills with SAS® Viya®

Jack McGuire, Amadeus Software Ltd.

ABSTRACT

Released in 2016, SAS Viya is the newest evolution in the SAS Platform, offering the latest high-performance analytics. SAS Viya brings a fundamental change to the SAS Platform, which has led to changes in how the platform is administered. As an experienced SAS9 administrator, the thought of having to deal with these changes seemed daunting, especially since learning that there is no metadata or SAS Management Console in SAS Viya. Fortunately, SAS Viya has been designed with the administrator in mind. A single web administration user interface, SAS Environment Manager, an HTML5 web application, centralizes the administration tasks.

This paper provides examples of common administration tasks in SAS Viya by relating them to the corresponding tasks in SAS 9. By comparing these tasks, administrators familiar with SAS 9 will not only gain confidence with SAS Viya but will also see the benefits that SAS Viya can offer them.

The intended audience for this paper is SAS 9 administrators who are interested in learning about SAS Viya 3.4 on both Microsoft Windows and Linux operating systems.

INTRODUCTION

With the introduction of the new SAS Viya platform, comes the need for a new set for skills for SAS Administrators. This paper aims to provide SAS administrators familiar with the SAS 9 platform the confidence to grow their skills and learn to administer a SAS Viya platform. To achieve this, we first look at the changes in architecture in SAS Viya before identifying the routine tasks a SAS Viya administrator is responsible for by comparing them to familiar tasks in SAS 9.

Later we investigate the SAS Viya routine administration tasks, which can be grouped into four key areas:

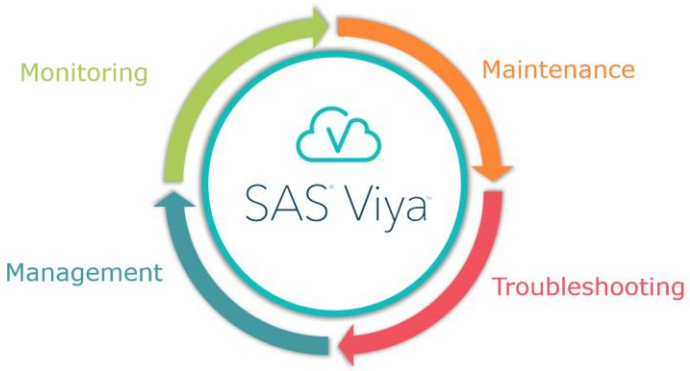

Figure 1: SAS Administrators Responsibilities

- Monitoring
  Performing systematic checks of platform availability and performance.
- Maintenance
  Carrying out regular periodic tasks to keep the platform secure, resilient and up to date.
- Troubleshooting
  Identifying and resolving any issues on the platform promptly.
- Management
  Maintaining processes to control user access and system backups.

SAS VIYA ARCHITECTURE

An understanding of the SAS Viya platform compared to the SAS 9 platform is vital for successful administration. This section looks at the architecture of the SAS Viya platform by breaking it down to its four main components, explaining the purpose of each, and drawing parallels to the SAS 9 platform.

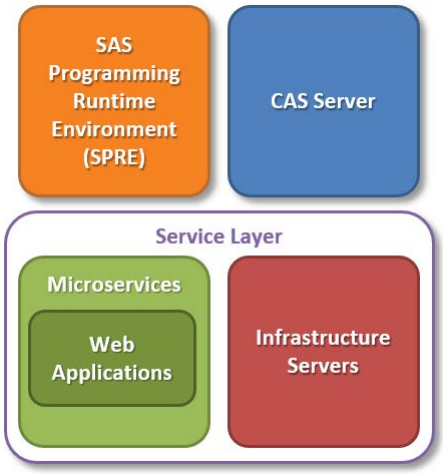

Figure 2: SAS Viya Architecture Components

SAS PROGRAMMING RUNTIME ENVIRONMENT (SPRE)

The SAS Programming Runtime Environment (SPRE) is the programming interface to SAS Viya which is accessed through SAS Studio. The SPRE should be the most familiar component to SAS 9 administrators as it is the equivalent of SAS Foundation on a SAS 9 platform.

The major benefit of the SPRE is that you can run legacy Base SAS code on SAS Viya without having to rewrite it.

There are two versions included:

- A legacy SAS 9 version, which uses the familiar SAS/CONNECT, SAS Object Spawner and SAS Workspace Server.
- An updated version, which utilizes the SAS Compute Server and SAS Launcher Server.

The SPRE component also allows SAS 9 clients like SAS Enterprise Guide to connect and run Base SAS code on the SAS Viya platform.

SAS CLOUD ANALYTIC SERVICES (CAS)

The SAS Cloud Analytics Services (CAS) is the latest SAS in-memory run-time environment. Deployed on a single server, the CAS server offers Symmetric Multi-Processing (SMP) which distributes the load across multiple processors. The full benefit of the CAS server is realized when it is distributed across multiple machines, this allows for Massively Parallel Processing (MPP), where the load is spread across multiple worker nodes. The CAS Server offers several benefits for high-performance analytics and distributed computing including dynamic scalability, virtual memory footprint, and parallel data loading.

SAS MICROSERVICES

The key change to the architecture is the move to use SAS microservices. Microservices are self-contained processes responsible for single aspects of the functionality of the software. SAS 9 platforms use Windows Services or Linux Daemons responsible for multiple functions and are interdependent. On SAS 9 platforms this means that the Windows Services or Linux Daemons must be restarted in a specific order; if you need to restart one component, you must also restart all the dependent ones too. With SAS Viya you can manage (start, stop and update) the microservices individually, reducing the downtime of the platform.

INFRASTRUCTURE SERVICES

The final component of the SAS Viya platform is the SAS Infrastructure services which are stateful processes, meaning that their state persists in some form of back-end storage and will do after a restart. These processes enable the microservices to remain stateless by providing persistent storage outside of the microservice. Without these stateful processes, it would not be possible to manage the microservices independently.

The infrastructure services are also responsible for providing components such as the web server to make the web interfaces available to the end-users.

VIYA ADMINISTRATION TASKS

After learning some of the key changes introduced in the Viya architecture, the next sections cover the routine administration tasks to maintain the platform.

MONITORING

In this section, we look at how you can monitor the SAS Viya platform and proactively identify and resolve issues early to ensure the best experience for the users.

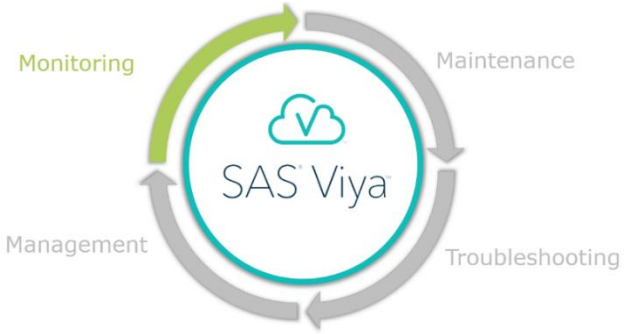

To make certain your organization gains maximum value from SAS Viya, you must ensure that the SAS Viya platform remains performant.

This section covers the performance monitoring of the SAS Viya platform focusing on the benefits of using SAS Environment Manager.

On a SAS 9.4 platform, SAS Environment Manager is a tool for platform performance monitoring, and you can configure it to extend the scope of the monitoring to give further visibility of the SAS environment. In SAS Viya, this monitoring tool has been improved extensively to display metrics in predefined SAS Visual Analytics reports (Figure 3: SAS Viya Environment Manager Monitoring Reports) which are accessible through the SAS Environment Manager Dashboard. The reports include metrics that are used to monitor the SAS Viya platform usage and performance including key metrics such as:

- Memory usage
- CPU usage
- Disk Space
- Network I/O usage

Using the SAS Visual Analytics reports allows you to drill down into individual user sessions to explore the performance of specific user processes.

Figure 3: SAS Viya Environment Manager Monitoring Reports

One example of this is using the 'Disk Space' report to monitor the key storage locations on the SAS Viya servers, such as the CAS backing store. These reports not only provide an overview of the current storage utilization but also provide a forecast to allow the administrators to predict when storage volumes are likely to fill up and take proactive steps to avoid this happening. The reports can also help with identifying the largest and fastest-growing directories by displaying the size of the directories over time.

Another example is the 'CAS Activity' report, which contains detailed performance monitoring metrics from the CAS servers. This report will help administrators to diagnose performance issues and ensure the server resources are sufficient and are being managed correctly. Administrators can see the current load on the CAS servers and drill down into metrics for the CPU, memory, I/O and threads for each CAS server (Figure 4: SAS Viya CAS Activity - CPU Usage Report).

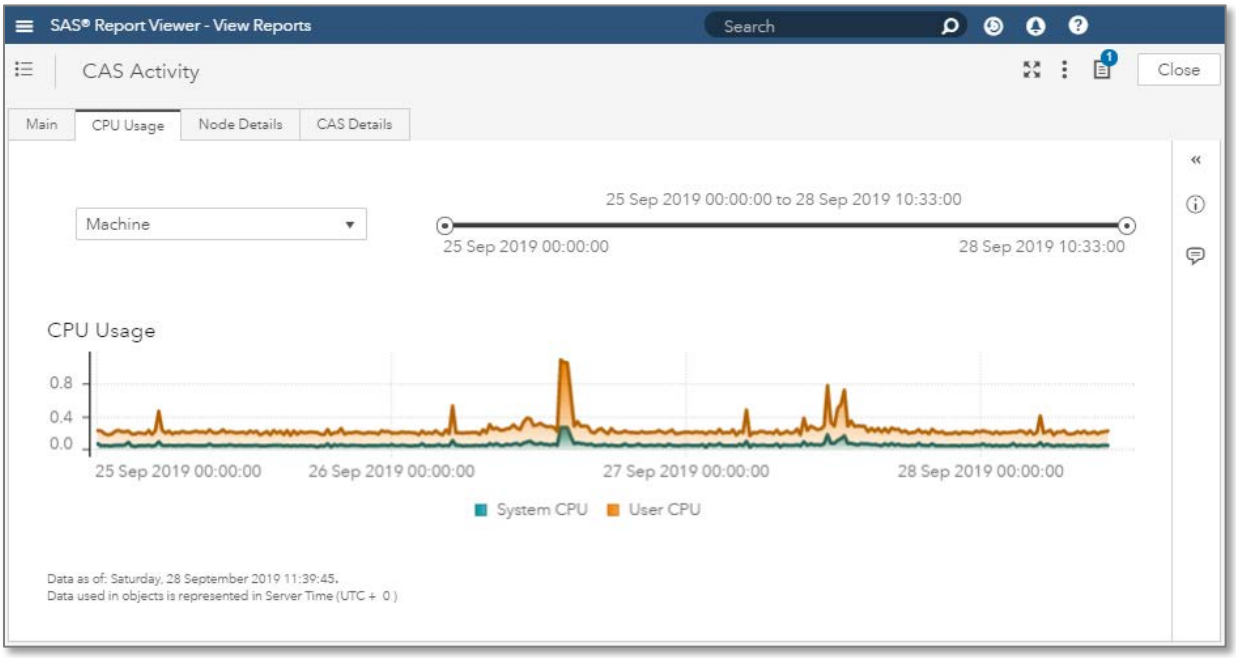

Figure 4: SAS Viya CAS Activity - CPU Usage Report

The data used to populate these reports is collected in tab-delimited (TSV) files. By default, an automated process loads this data into CAS, compresses the TSV files into zip files and eventually deletes them. You can view the properties of this process by running the following command on the SAS Viya server:

```
ops-dm-admin show
```

The table below shows some of the configurable properties:

Property	Description	Default Value (days)
EMI_CAS_RETAIN_DAYS	Number of days the data is retained in the CAS server	3
EMI_DELETE_TSVZIP_DAYS	Number of days before data is deleted from the data mart	10
EMI_ZIP_TSV_DAYS	Number of days before the TSV files are compressed into zip files	1

Table 1: SAS Datamart Properties

These properties can easily be modified by using the following syntax:

```
ops-dm-admin set property=value
```

For example, to retain five days of data in the CAS server instead of the default three, you can run the following command:

```
ops-dm-admin set EMI_CAS_RETAIN_DAYS=5
```

```
[sas@viya  bin]$ ./ops-dm-admin set EMI_CAS_RETAIN_DAYS=5
2019-09-29T11:07:46.531506+01:00 INFO  [main.go:255] [ops-dm-admin] - Datamart set
        config/ops/datamarts/evdm/ARM_SUBSYS=ARM_PROC
        config/ops/datamarts/evdm/EMI_CAS_LOAD=Y
        config/ops/datamarts/evdm/EMI_CAS_LOAD_LOGS=Y
        config/ops/datamarts/evdm/EMI_CAS_LOAD_METRICS=Y
        config/ops/datamarts/evdm/EMI_CAS_RETAIN_DAYS=5
        config/ops/datamarts/evdm/EMI_DELETE_TSVZIP_DAYS=10
        config/ops/datamarts/evdm/EMI_ZIP_TSV_DAYS=1
```

Figure 5: Modified Datamart Values

MAINTENANCE

In this section, we look at regular periodic tasks that an administrator should carry out to ensure the platform is kept secure, resilient and up to date.

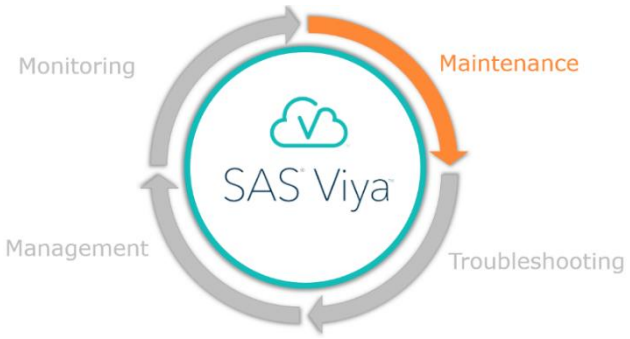

UPDATES

Ensuring the SAS software is up to date, with the latest security patches or bug fixes, is an important administration task. In SAS 9 we apply updates as hot fixes, and using the SAS Hot Fix Analysis, Download and Deployment Tool (SASHFADD) we can identify any available hot fixes and apply them using the SAS Deployment Manager. The complete process involves several steps in order to identify, download and ultimately apply relevant hot fixes (Figure 6: SAS 9 Hot Fix Process).

Figure 6: SAS 9 Hot Fix Process

In SAS Viya updates have been streamlined and usually involve a single operating system command, similar to the command used to deploy the software (Figure 7: Example Update Command (Single Server Ansible Linux Deployment)). This command checks the installed software against the latest repository made available by SAS before downloading and applying the relevant patches in a single step.

```
[root@viya  sas]# ansible-playbook site.yml
```

Figure 7: Example Update Command (Single Server Ansible Linux Deployment)

CLEANWORK

In SAS Viya, you can run existing SAS 9 code without any modifications. In order to facilitate this process, a SAS Work library is still required by applications such as SAS Studio. As with SAS 9 platforms, there is a risk that temporary files can be left behind if the session does not end cleanly. To address this problem, the SAS Cleanwork utility is still available for the SAS Viya platform and we should configure and maintain it in the same way as for a SAS 9 platform.

LOG ARCHIVING

On SAS 9 platforms, a routine housekeeping task is to remove or compress old log files. These log files are stored in many directories throughout the configuration directory, complicating this process. In SAS Viya the log files are configured to be stored in a single directory. The SAS Viya platform also has a default process to archive log files older than 30 days and compress them into daily zip files, reducing their size by around 95%.

Although this process does help reduce the amount of disk space taken up with log files, it is still important to implement a process to manage these zip files. The zip files can be deleted or moved to another storage device, depending on your organization's retention policies.

To simplify this task, all the daily archived log files are stored in the same directory as a .zip files as shown in the screenshot below (Figure 8: SAS Viya Log Directory).

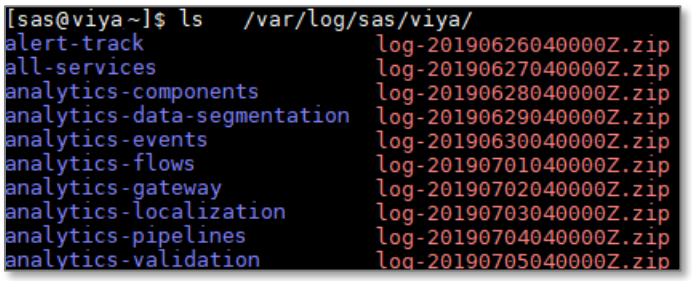

Figure 8: SAS Viya Log Directory

LICENSING

SAS Viya, like SAS 9, has a license file that includes the specific products your site has licensed and an associated expiry date for each. Administrators must update this license file before the expiration date to avoid disruption.

As with SAS 9, we can check the status of the current license using the SETINIT procedure through SAS Studio. In SAS Viya we can monitor the status of the license using SAS Environment Manager too, which now includes a page displaying the following information on the licensed products (Figure 9: SAS Viya Environment Manager: Licensed Products):

- Product Name
- Product ID
- Status
- Expiration Date
- Grace Period End
- Warning Period End
- Maximum CPU Count

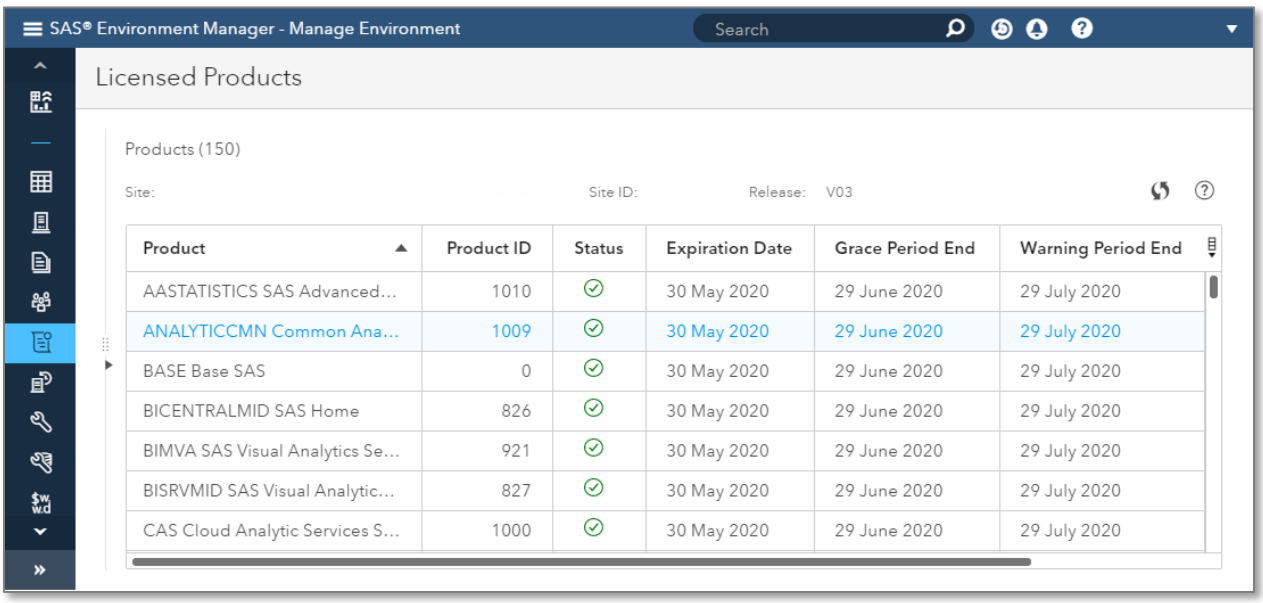

Figure 9: SAS Viya Environment Manager: Licensed Products

The process for updating the license varies from SAS 9 because we can complete it from a single server and the license is automatically distributed across the platform. We need to apply the new license from the deployment directory using Ansible (Linux) or PowerShell (Windows).

First, we update the `vars.yml` (Linux) or `vars.psd1` (Windows) to point to the new license file, then run the following command from the location where the deployment scripts are located:

`ansible-playbook apply-license.yml` (Linux)

`setup.bat -apply-license` (Windows)

After we have applied the license, any newly created CAS sessions will pick up the information from the new license file without restarting the platform, any existing sessions will continue to run unaffected using the previous license.

TROUBLESHOOTING

This section covers the new features in SAS Viya which assist SAS administrators in troubleshooting issues on the SAS Viya platform.

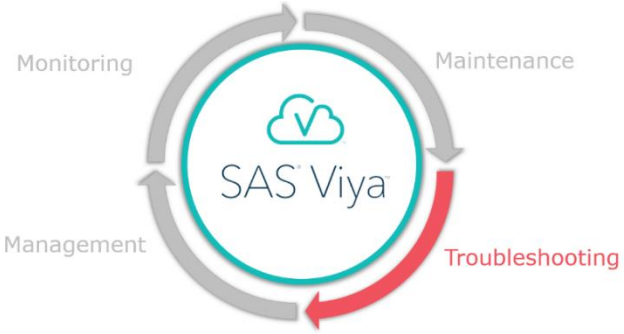

SERVER STATUS

To ensure the platform is up and running on a SAS 9 platform, the administrator can check that the Windows services or Linux Daemons are running. In SAS Viya there is an additional way to quickly check the status of the SAS Platform. On the home page of the SAS Environment Manager web application, there is a 'Dashboard' screen with an 'Availability' tile which shows the status of each component (Figure 10: SAS Viya SAS Environment Manager Dashboard). The 'Availability' tile contains three sections showing increasing levels of granularity:

- Machines – all the hosts where SAS Viya is running, clicking on a machine will highlight all the services and service instances running on that machine. The screenshot below shows a single-machine deployment.
- Services – all the running infrastructure services and microservices.
- Service Instances – all the processes running for the microservices and infrastructure services, some services have multiple processes associated with them.

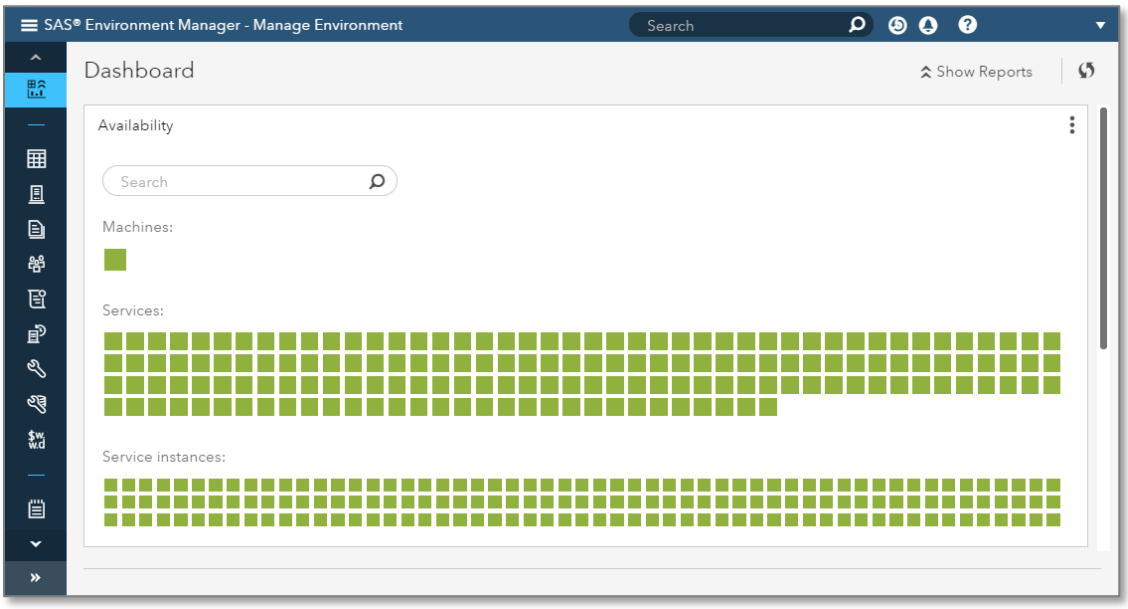

Figure 10: SAS Viya SAS Environment Manager Dashboard

The tile shows the status of each machine, service and service instances, with a color-coded box. If the component is available, the box is green, if there is an issue or warning the box is yellow and if the component is unavailable, the box is red. If the service is not running, it will not be displayed in the 'Availability' tile and must be started from the server. The search bar can be used to find a specific component. Hovering over one of the boxes will provide the name of the machine, service or service instance. Clicking on a service instance will highlight the related service and machine and provide information such as the port number and IP address of the host.

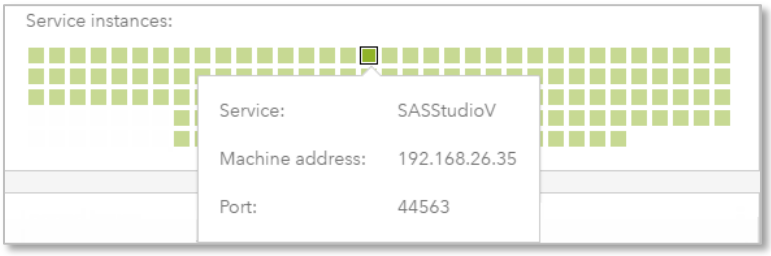

Figure 11: Service Instance Detail

We can still check the status of the services (both Infrastructure Services and Microservices) using Windows Services or Linux Daemon scripts, however, the SAS Environment Manager dashboard offers a centralized location where we can monitor them without requiring access to each of the SAS host machines individually. Aside from convenience, another benefit is that we do not need to connect to the server directly through remote desktop or/and an SSH connection.

LOG MANAGER

While troubleshooting a SAS 9 platform, it is common to access the server filesystem to open the most recent log files and search for errors or warnings. This usually involves connecting to each machine in the deployment and navigating around the SAS configuration directories for the different SAS log files. In SAS Viya, we can review the log messages centrally from SAS Environment Manager. The 'Logs' tab of SAS Environment Manager contains messages from all the SAS logs in a standard format (Figure 12: SAS Viya Environment Manager - Logs Tab).

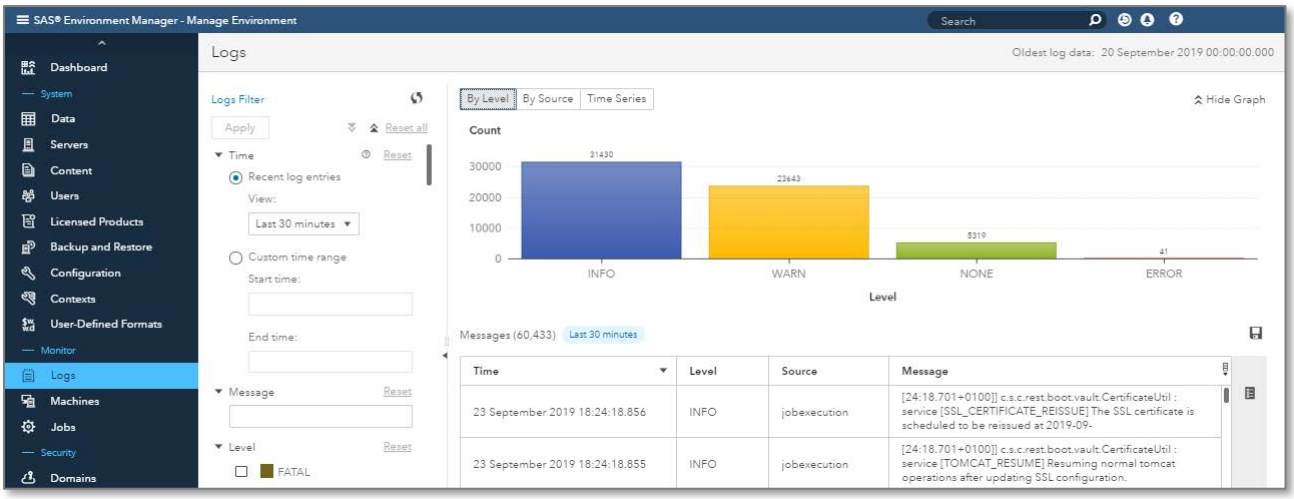

Figure 12: SAS Viya Environment Manager - Logs Tab

We can sort these log messages in chronological order and filter by message level, source, date and time. This clearly simplifies the process of identifying the first occurrence of an error, making it easier to follow the error chronologically through all the log files without having to open and close individual files or connect to multiple machines. The screenshot below (Figure 13: SAS Viya Environment Manager - Logs Tab Filter for Errors) shows the output when filtering for errors in the last 30 minutes from the source CAS.

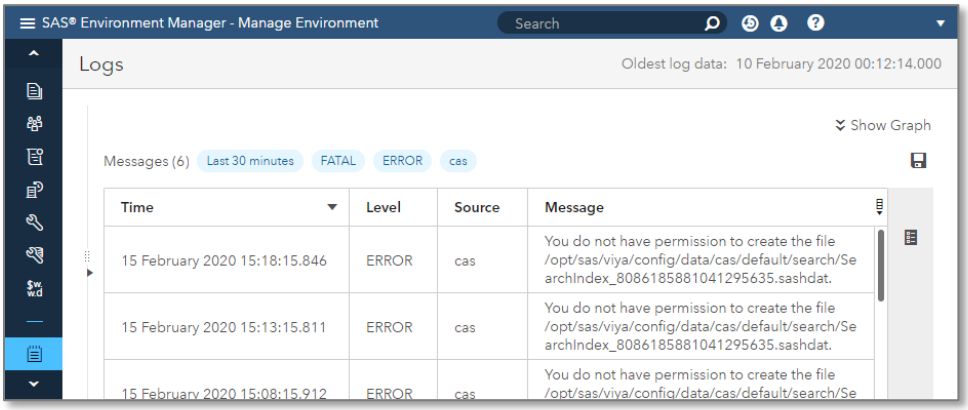

Figure 13: SAS Viya Environment Manager - Logs Tab Filter for Errors

MANAGEMENT

In this final section, we cover the processes to control user access and system backups.

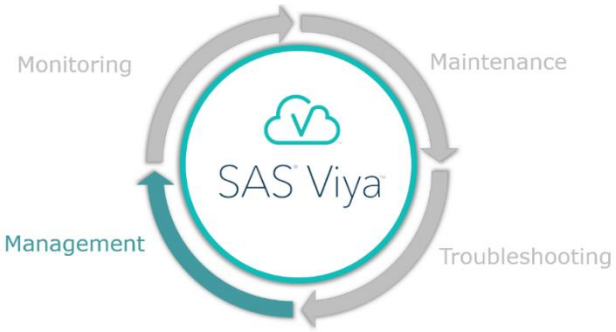

USER MANAGEMENT

A common request for a SAS administrator is to set a user up with access to SAS, in SAS 9 this involves creating a SAS metadata user account using SAS Management Console. In SAS Viya, there is no need to create an account manually as SAS Viya integrates directly with your identity provider, such as Active directory, to synchronize the user accounts. Once we have set up a connection, the users defined in the connection will automatically be available in the 'Users' tab of SAS Environment Manager.

Now the primary administration task is to ensure that we only synchronize relevant users by reviewing the connection details, which can be done through the 'Configuration' tab in SAS Environment Manager (Figure 14: SAS Environment Manager - Identities Configuration). The 'Identities Service' configuration contains three sections:

- LDAP Connection – Configuration of the connection to the identity provider.
- LDAP Group - Configuration of which groups to synchronize from the identity provider.
- LDAP User – Configuration of which users to synchronize from the identity provider.

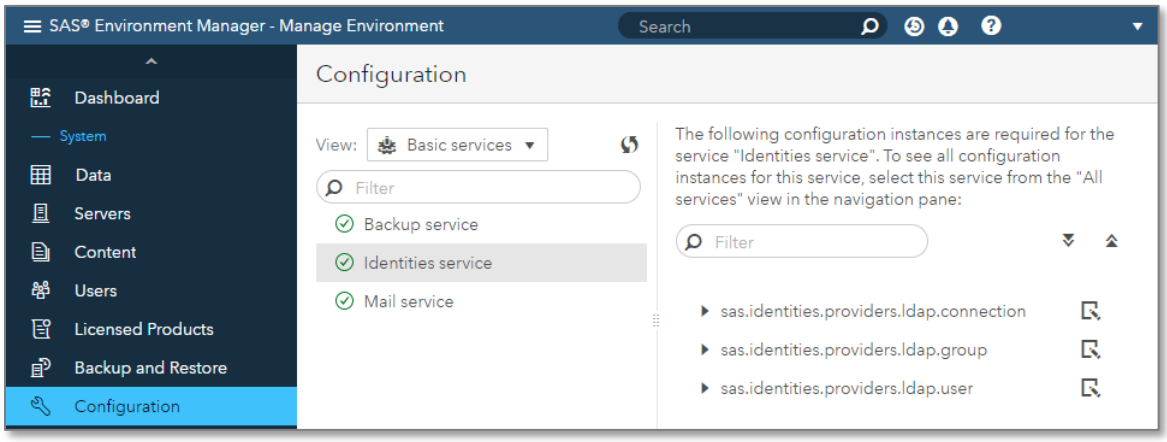

Figure 14: SAS Environment Manager - Identities Configuration

For both the LDAP User and LDAP Group configuration, the `baseDN` property tells SAS in which directory of your identity provider to find the user and groups. By refining this property, we can ensure to synchronize only relevant users and groups.

The value for the `baseDN` property takes the following form:

```
CN=Users,DC=Domain,DC=com

CN=Groups,DC=Domain,DC=com
```

Where the value of the `objectFilter` property further refines the synchronized users and groups.

One of the principal benefits of SAS metadata accounts in SAS 9 is the extra layer of security offered by SAS metadata roles and permissions. This extra layer of security is still possible in SAS Viya, however, instead of having to create groups manually in SAS metadata, SAS Viya can synchronize groups directly from your identity provider. If we require more granular permissions, we can create 'Custom Groups' in SAS Environment Manager to manage permissions and access.

SAS BACKUPS

As SAS administrators, we should ensure we are prepared for the worst-case scenario by having a defined process for backing up and restoring the SAS platform and its content. Then if issues do occur, we can be confident the content and platform can be restored quickly and safely minimizing disruption.

The release of SAS 9.4 (TS1M3) brought the much-welcomed SAS Backup Manager web application (in earlier versions of SAS 9.4, we managed the backups through a series of batch commands on the operating system). The SAS Backup Manager allowed administrators to take backups, manage backup schedules and backup configuration using a point and click interface through the web application (Figure 15: SAS 9.4 (M3 onwards) Backup Manager).

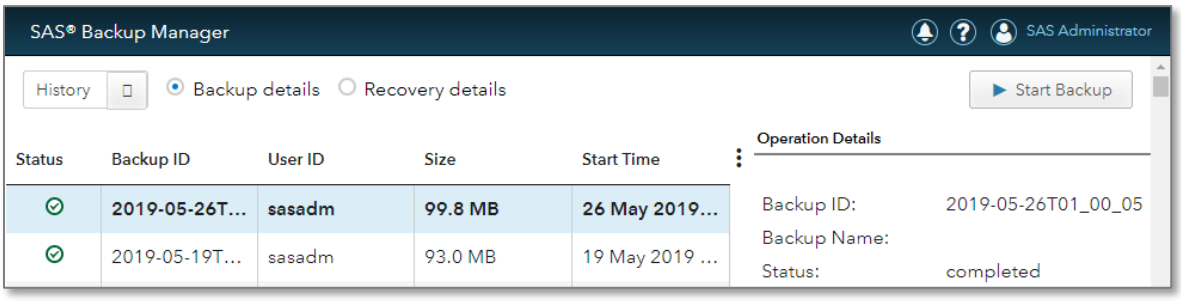

Figure 15: SAS 9.4 (M3 onwards) Backup Manager

In SAS Viya, the latest evolution of the SAS Backup Manager has a familiar user interface, although several additional features have been introduced to aid the SAS administrator. Most notable is the addition of the restoration of backups through the web application. In SAS 9.4 we completed this task using operating system batch commands but now a new 'Restore' action has been added to the Backup Manager (Figure 16: SAS Viya Backup Manager - Restore Feature).

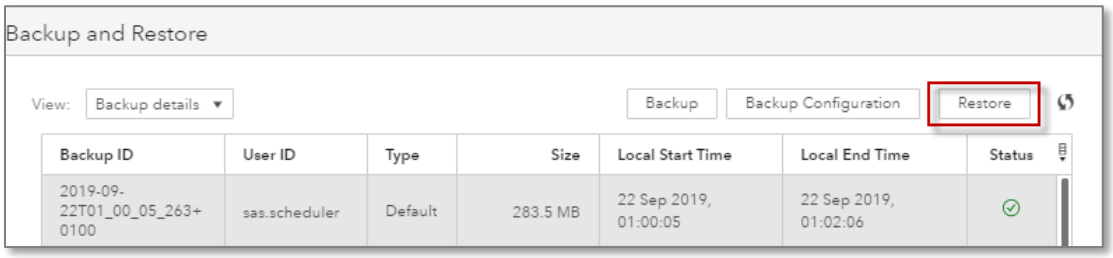

Figure 16: SAS Viya Backup Manager - Restore Feature

There are other components that are vital for a successful backup strategy. These elements sit outside the SAS Backup and Recovery took and they should be considered as part of the wider backup strategy:

- Operating environment

- SAS Viya deployment – for example the installation and configuration directories
- User home directories
- Data sources – for example, the data in CAS, databases and on local and remote file systems.

In practice, if disaster recovery is required, virtual machine snapshots or checkpoints offer the most complete backups and quickest recovery option for most platforms. Any backups should be executed at regular intervals and it is recommended that these backups are stored on a separate device from the SAS environment.

Any organization should have a backup and recovery strategy clearly defined and tested at periodic intervals to provide confidence that in the event of a disaster the platform can be recovered quickly and safely.

CONCLUSION

In this paper, we have seen how the SAS Viya platform differs from the SAS 9 platform architecturally and the benefits this offers to organizations. For example, the introduction of microservices, which have fewer interdependencies, enable administrators to restart individual functions and therefore reducing downtime for the end-users. We have covered how monitoring platform availability and performance is smoother with the predefined and interactive SAS Visual Analytics reports. Many of the periodic maintenance tasks are familiar in that they achieve the same outcome, but the process has been streamlined. When things go wrong, the troubleshooting steps are essentially the same, however, using the 'Logs' tab in SAS Environment Manager helps to speed up the process. Finally, we looked at management tasks, such as managing users which now contains less manual input from administrators with the introduction of synchronization from your identity provider.

Although at first sight, administering a SAS Viya platform for the first time may seem intimidating, the tasks covered in this paper are not dissimilar to the SAS 9 administration tasks and more often than not they have been simplified and improved, making the administrator's job more efficient. The SAS Viya platform offers SAS administrators an excellent opportunity for continuing to develop their existing administration skills.

REFERENCES

SAS Institute Inc. 2018. SAS Viya 3.4 Administration. Cary, NC: SAS Institute Inc. Available at https://documentation.sas.com/?cdcId=calcdc&cdcVersion=3.4&docsetId=calchkcfg&docsetTarget=n00004saschecklist0000config.htm&locale=en

RECOMMENDED READING

David Stern (2019) Checklist of SAS Viya Administration Tasks: SAS Institute Inc.

CONTACT INFORMATION

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

Jack McGuire
Amadeus Software Ltd.
+44 (0) 1993 848010
jack.mcguire@amadeus.co.uk
www.amadeus.co.uk