

Paper SAS4313-2020**How to Maintain Happy SAS® Viya Users**

Margaret Crevar, SAS Institute Inc.

ABSTRACT

Today's SAS® Viya® environments support many concurrent processes, using ever-growing data volumes. To help SAS Viya users remain productive, SAS® administrators must ensure that the SAS Viya applications have both sufficient and properly configured compute resources that are continuously monitored. Understanding how all the SAS components work and how they are employed by your users is the first step. The guidance offered in this paper helps SAS administrators configure and tune SAS Viya hardware, operating systems, and infrastructure. This tuning and configuration will keep their SAS Viya applications running at optimal performance and keep their user community happy.

INTRODUCTION

There is one common thread we have found regarding maintaining happy SAS Viya users when working with IT administrators at customer sites. That thread consists of ensuring that the underlying hardware is properly configured to support the SAS Viya applications. These installations have reported the best SAS Viya user experiences. Obtaining optimal hardware configurations is not a trivial task. Different SAS Viya technologies often require hardware configurations that vary from the traditional SAS®9 or relational database environments. Understanding how SAS Viya is used at your site helps you evaluate options for the hardware, operating system, and infrastructure (middle tier) configurations. This is easier for existing SAS®9 customers than for new SAS Viya customers or new SAS Viya applications at an existing SAS9 customer site.

In this paper we will do the following:

- discuss briefly how SAS Viya works, focusing on the communications between the components
- discuss how SAS Viya can be used
- discuss how SAS Viya is different from SAS9
- offer guidance on how to initially configure hardware for SAS Viya usage
- offer guidance on how to use virtualization with SAS Viya servers
- provide advice on how to monitor the hardware to avoid running out of a computer resource

This paper consolidates information that has been presented in previous SAS Global Forum papers. The goal for consolidating the information into a single paper is to help the SAS and IT administrators who must support SAS Viya better understand how SAS Viya works and what they should do to keep their SAS Viya customers happy.

HOW SAS VIYA WORKS

SAS Viya is our latest extension of the SAS® Platform. Designed to enable analytics for an enterprise, it scales data operations of any size, type, speed, and complexity, and is

interoperable with SAS 9.4. In addition to individual offerings and solutions, SAS Viya is also available in five analytical life cycle bundles.

SAS Viya users have access to both the distributed computing capabilities of the SAS Viya products and the functionality of the included SAS 9.4 products. You can run your new SAS Viya code as well as your existing SAS code in your SAS Viya environment, which can access both the SAS Foundation server and the CAS server. Please refer to the technical architecture diagram in Figure 1.

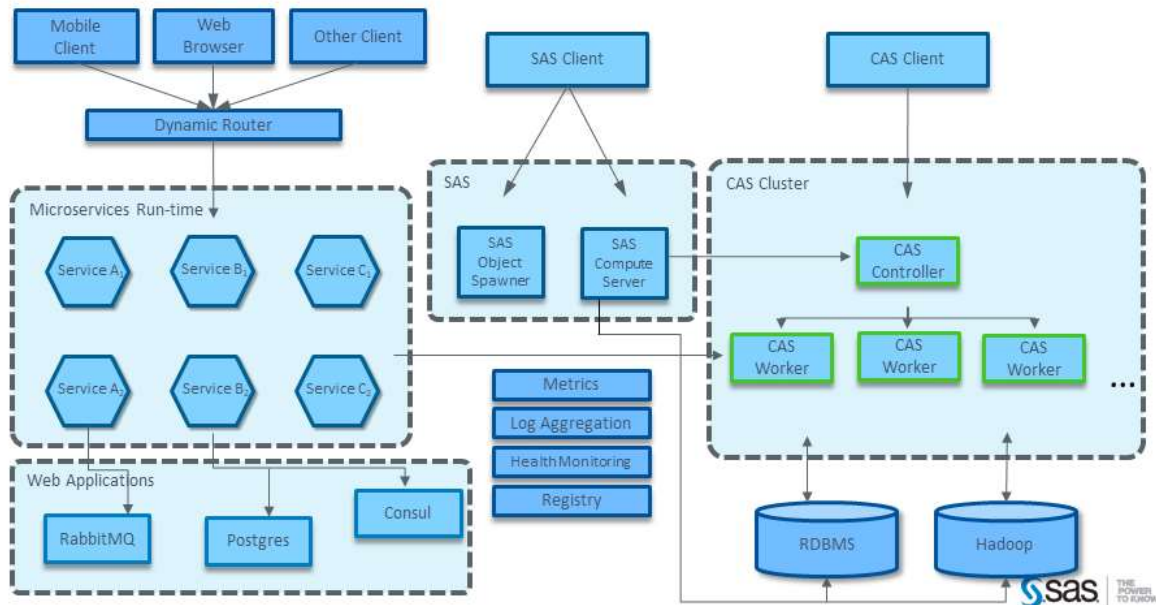


Figure 1: SAS Viya Architecture

If you are a SAS 9.4 customer who is looking to move your existing SAS9 applications to SAS Viya, you need to understand the optimal places within SAS Viya to execute these applications. The two places in question are in the SAS compute server or in the CAS environment. For more information about optimal task and job placement, please refer to "Comparing SAS Viya and SAS 9.4" (Schneider 2019). Please note that Mark is giving a paper at SAS Global Forum 2020 entitled "Migrating from SAS 9.4 to SAS Viya".

UNDERSTANDING SAS VIYA USAGE

SAS Viya usage can span from heavy analytics to simple query and reporting. The IT administrative staff needs to understand how SAS Viya will be used and how SAS Viya interfaces with the computer hardware, networks, and storage it is running on. The better their understanding of how SAS Viya works, the better they can configure the computer resources. This is especially true for the I/O subsystems, file systems, and memory that supports the SAS Viya users at their site. The SAS Viya users need to work with their IT administrators to review and characterize their tasks and interactions. The IT administrators can then use the general guidelines in this paper to set up a Linux system to support SAS Viya jobs.

The following topics discuss different types of SAS Viya applications. Please note that, although there are many individual SAS Viya applications, the applications that are discussed in this paper tend to be the most widely used.

SAS® VISUAL ANALYTICS

SAS Visual Analytics is an interactive reporting and visual discovery product. It is designed to keep all the data needed for reports in memory. Typically, many SAS Visual Analytics users share a single table for their reporting. The majority of users tend to perform lightweight analytics with lesser amounts of I/O. This tends to be the highest SAS Viya usage category.

SAS® VISUAL STATISTICS, SAS® VISUAL FORECASTING, SAS® VISUAL TEXT ANALYTICS, AND SAS® VISUAL DATA MINING AND MACHINE LEARNING

SAS Viya offers complex analysis and modeling applications such as SAS Visual Statistics, SAS Visual Forecasting, SAS Visual Text Analytics and SAS Visual Data Mining and Machine Learning. SAS Viya users performing tasks associated with these applications tend to use larger, non-shared data tables, and the analytical tasks might create additional copies of the data during their execution. The heavier analytical algorithms and data manipulation associated with these applications tend to employ significant I/O, both in memory (as CAS memory mapped files in virtual memory), and on device (the on-device portion of the CAS memory map known as the CAS_Disk_Cache. There tend to be fewer concurrent analytics sessions with these tasks as compared to SAS Visual Analytics.

EXTRACT, TRANSFORM, AND LOAD (ETL) PROCESSING

As with any SAS application, there is often the need for extracting data, transforming it, and loading it into tables before analysis can be performed. With SAS Viya, you need to make the decision if this ETL processing will be done in CAS or on the SAS compute tier. In either case, the need for extra memory (if done in CAS) or disk space (if done in SAS) need to be taken into account. In addition, you need to have a complete understanding of the ETL process cycles, in order to avoid resource contention with SAS Viya Application users in the same time period.

GUIDANCE ON CONFIGURING HARDWARE RESOURCES FOR SAS VIYA USAGE

Now that we have a basic understanding of SAS Viya, let's discuss what hardware resources are needed. Configuring hardware to support the profile of your SAS Viya applications makes a significant impact on overall performance. When your computer resource must support multiple SAS Viya applications (especially if their access patterns vary widely), you might need to make tradeoffs (performance versus price is the key one).

MEMORY:

When you use CAS, the tables that are used are loaded into memory as memory-mapped files, so you need enough physical RAM to support all the tables currently being used, including any tables or working data created during your session. If you do not have enough physical memory, then CAS might be forced to page portions of a table out. If they are subsequently required, they will have to be restored from the backing store (original source or CAS_DISK_CACHE – the on-device portion of the memory map depending on the table type and location). Please note that CAS_DISK_CACHE is on physical devices (HDD, SSD, NVME), so the performance to retrieve pages from these devices will be significantly slower than in-memory operations. More details on configuring CAS_DISK_CACHE are provided in the I/O section. These are guidelines for the memory required:

- CAS controller and workers: Enough space to hold all tables and partitions expected for the concurrent workload. There should be at least 8 GB of RAM per physical core
- microservices tier: A minimum 8 GB RAM per physical core
- web applications tier: A minimum of 8 GB of RAM per physical core

- SAS compute tier: A minimum of 8 GB of RAM per physical core

NETWORK:

When CAS actions and procedures execute, a lot of communication takes place between the SAS Viya nodes (such as microservices, RabbitMQ, Postgres, and so on). In addition, large-block data that the nodes process is also carried on the same pathways that are used for these communications. Because of this mix of large data shipping and significant small services communications, we recommend a minimum of a dedicated 10-gigabit network connection for all nodes. Not having a robust enough network between the SAS Viya nodes is one of the leading reasons for instability and performance issues with SAS Viya deployments.

TIP: If your system has multiple NIC cards, please read this SAS Communities [blog](#) "Understanding Network Binding Variables in SAS Viya 3.5" (RobCollum, 2020)

CPU:

The analytical tasks described previously for the SAS Viya applications tend to be more CPU-intensive in nature. These job tasks within these applications are designed to start a thread for every physical core on your CAS workers. Therefore, it is not uncommon for a single job to consume an entire set of processors within your system until it is finished with its analytical processing.

TIP: To maintain happy SAS Viya users, you need to both determine how many of these long-running analytical tasks will be happening simultaneously and make sure you have enough CPU cycles to support them along with your other SAS users' tasks.

TIP: It is strongly advised that the CAS workers should be homogenous systems with the same chipset, amount of RAM and number of physical processors. SAS Viya tasks work in parallel, so it is very helpful if each parallel thread runs on consistent resources and execution time of each parallel task can be equalized across the threads and nodes.

TIP: The CAS controller should be the same chipset as the CAS workers, but it can have fewer physical processors and RAM.

I/O:

There are two major file systems that can influence the speed of the CAS workers. The first is the persistent data file storage. The I/O throughput to this file system will influence the amount of time it takes to load data into memory for processing. The second is the CAS_Disk_Cache (CDC) file system that is used as an on-device portion of the memory map. Files that require backing stores will be housed in the CDC, or in its original backing store source. If an in-memory portion of file is paged out, and subsequently needed, it will go to the backing store to retrieve it. For files that follow rules to use the CDC as a backing store, they file portions will be retrieved from there. It is a backing store for availability and safety when file pages are needed by CAS PROCEDURES and actions but are no longer in virtual memory.

TIP: Move the pointer for CAS_DISK_CACHE from the /tmp directory to a more robust file system.

TIP: Some additional reading on how to configure CAS_DISK_CASCHE can be found here:

- Information on how to properly configure CAS_DISK_CACHE can be found in the SAS Global Forum paper "Engineering CAS Performance Hardware Network, and Storage Considerations for CAS Servers" (Brown 2019).

- SAS Communities Administration [blog](#) "Provisioning CAS_DISK_CACHE for SAS Viya" (RobCollum, 2019)
- SAS Communities Administration [blog](#) "4 Rules to Understand CAS Management of In-Memory Data" (RobCollum, 2019)

TIP: In order to avoid I/O conflicts, the system administrator or installer should be cautious sharing the underlying physical storage of these heavily used I/O file systems with other applications (whether these applications are performing heavy I/O or randomly accessing the data). If the storage must be shared, as in the case of "striped everything" spinning storage, you must carefully monitor and provision the storage. If possible, it is wise to not share highly random I/O with underlying SAS provisioned resources.

TIP: CAS_DISK_CACHE cannot be placed on IBM Spectrum Scale or Red Hat GFS2 file systems.

The caslib that is used by the SAS Viya users determines which of these file systems are most critical in an MPP CAS infrastructure.

- Path: Disk I/O for file acquisition is performed by the CAS controller and distributed to CAS workers in MPP systems.
- DNFS: Disk I/O for file acquisition is performed by the CAS worker nodes in parallel.
- CAS_DISK_CACHE:
 - Data blocks that are changed in virtual memory by CAS PROCEDURES or actions, or new data blocks created are backed in the CDC. Make sure you have enough device space for the SAS Viya analytical tasks that create extra data (e.g. sums of squares tables, model trees, etc.) when analyzing the data
 - SAS Model Studio recommends a minimum of 3 times the size of the data source.
 - SAS Visual Forecasting recommends a minimum of 10 times the size of the data source for early versions of SAS Viya. The requirement has been reduced to 3 times the size of the data source with the 19w21 release of SAS Viya.
- CASDATADIR: Move it to a high-performance storage platform with high IO bandwidth, especially when using SAS Visual Forecasting, SAS Visual Data Mining and Machine Learning and SAS Visual Text Analytics.

Figure 2 illustrates the major I/O methods discussed in this topic in order to make it easier for you to understand how I/O travels between the many SAS nodes.

Serial SAS9 vs. Parallel CAS

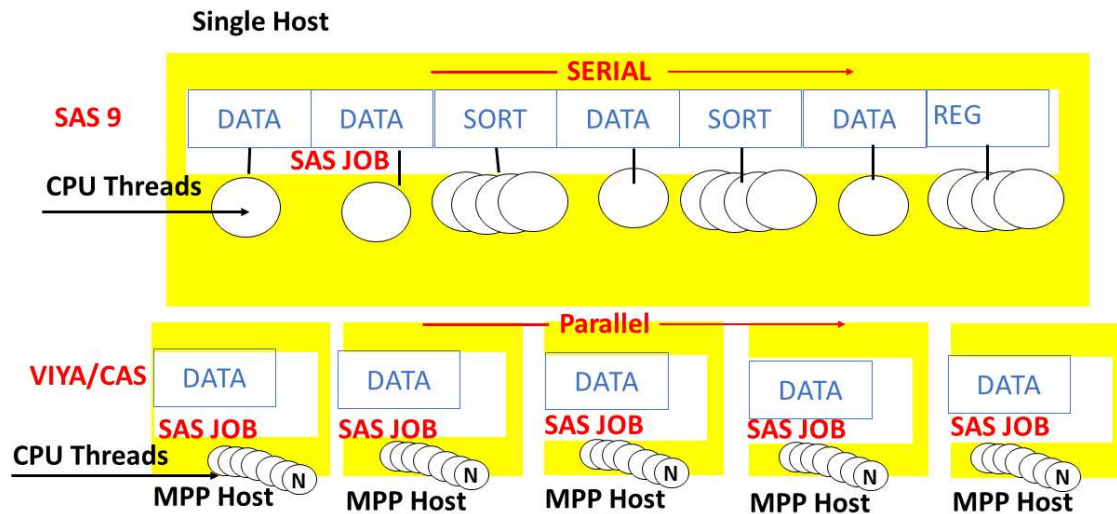


Figure 2: I/O in SAS Viya and SAS®9

In addition to the CAS controller and workers, we must consider there is the I/O throughput needs of the SAS compute server. This is the server that performs SAS 9.4 tasks. The I/O throughput requirements are the same as for any SAS 9.4 system. Consult with your storage administrator regarding how to provide your throughput needs based on these general guidelines:

- Overall I/O throughput needs to be a minimum of 100–125 MB/sec/core.
- SASWORK needs a minimum of 150 MB/sec/core.
- Permanent SAS data files need a minimum of 75–100 MB/sec/core.

Ensure that the heaviest-used file systems (usually SASWORK, permanent storage, UTILLOC, and so on) can sustain the I/O throughput requirements in the previous list. Again, consult with your storage administrator for help.

TIP: Move the file system used for SASWORK from the `/tmp` directory (the default location) to a more robust file system that meets the I/O throughput recommendations.

The initial disk space allocated for each of these file systems depends on the SAS application. An ETL job that transforms data from one source to a SAS data warehouse format generally has many steps that create temporary files, so you need to make sure that you have enough disk space for all the temporary files that are created. The same is true with ad hoc jobs that might be run by your end users. However, some BI applications or SAS solutions might not create any temporary files.

- **A large number of temporary files can be created during long-running SAS jobs.** These files are not only created, but might also be renamed, deleted, and manipulated many times during a long-running SAS job. The size of the files might range from very small (under 1 GB) to very large (in the 100s of GBs). The location of these files is referred to as the SASWORK area.

TIP: You must ensure that SASWORK on the SAS compute server points to a file system that not only has enough I/O throughput to support your users' I/O demands but also enough disk space to support all of the temporary files

that are created during each SAS session. This file system needs to support 50% writes, 50% reads, and 100% destruction.

A good starting point for the size of system resources are:

- CAS_DISK_CACHE is two times the size of the physical RAM.
- SASWORK is best sized by using this formula:
(number of concurrent steps in the concurrent running jobs) X (size of the input file)
X (number of current SAS jobs running).

For a better value for the number of file systems needed and their size, a detailed technical assessment is required. Please work with your SAS Account Team on how to accomplish this.

See the SAS Global Forum paper "Best Practices for Configuring your IO Subsystem for SAS®9 Applications" for more detail about the RAID levels and preferred tools to use to set up the file systems for the SAS compute server.

GUIDANCE IN TUNING SAS VIYA OPTIONS

Through helping SAS Viya users, the past few years, we have uncovered some SAS Viya parameters that you might want to tune from their default value:

- Set the `cas.memorysize` option to 80% of RAM on the CAS workers.
- If you are not using HDFS caslibs, set the `cas.colocation` option to "none". This requires the `cas.colocation='none'` statement be added to the `/opt/sas/viya/config/etc/cas/default/casconfig_usermods.lua` file and a restart of the server.
- For the `_readAhead` parameter, change the default from "false" to "true" with the CAS action `table.loadTable` when used with DNFS reads of large files on some file systems and file types. When loading large SASHDAT files via DNFS, changing this option from "false" to "true" changed the time required to initially load non-captured pages from 29 minutes to six seconds. When the value is "false", the initial read only reads section headers, not each page of the file. Subsequent re-reads from the source to capture pages beneath the headers are expensive. Some file systems and file types see the multi-thread pattern as a skip sequential/random operation and yield poor transfer performance. We have found this true in IBM Spectrum Scale (GPFS) and S3, and we are testing other systems.
- For the `_casout=(columnar=true)` parameter, you should trigger parquet output.
- Set the default value of the `copies/replicas` parameter to 0 (no replicas). This prevents CAS worker nodes from backing their CDC blocks to other nodes for failover. This does impute some risk for high availability (HA), but not replicating the data can alleviate network and host pressure on very busy systems. Utilizing a device stripe >R0 vs. JBODs might provide alternate protection.
- For SAS Visual Analytics users, you might consider reducing the number of CAS sessions that are automatically started for SAS Visual Analytics sessions from 6 to 3. Follow these steps:
 1. In SAS Environment Manager, select **Configuration** -> **All Services** -> **Report Data service**.
 2. Edit the configuration for the `sas.reportdata.properties`.

3. Scroll to the bottom and select the **Add property** option. Add these properties:
 - pooledCasSessionMaximum = 3
 - pooledCasSessionMaximumIdle = 3
4. Restart the report-data service to cause the change to take effect. Please keep the values in sync.

In addition to specifying these SAS Viya parameter settings, review the SAS Global Forum paper “Seriously Serial or Perfectly Parallel Data Transfer with SAS® Viya” (Collum 2019) that discusses how to load data quickly into CAS.

GUIDANCE ON TUNING HARDWARE RESOURCES

Now that we have discussed how SAS Viya works in general and have examined how to perform simple sizing for the hardware, we can look at some general configuration tuning issues to consider for the hardware where you run your SAS® Viya applications.

- **Red Hat Enterprise Linux (RHEL) tuning:** Follow the best practices in the Red Hat publication “Optimizing SAS on Red Hat Enterprise Linux (RHEL) 6 & 7”.

In addition, you should tune these RHEL kernel settings:

- vm.min_free_kbytes = ~1GB
 - vm.dirty_background_ratio = 3
 - vm.dirty_ratio = 10
 - RHEL turn huge transparent page off
 - RHEL turn memory defragmentation off
 - Multi-NIC cards best practices – point to
 - Slow NIC – DNS
 - Fast NIC – internal communications
- **ulimit Values:** Follow the guidelines in the “SAS Viya 3.5 for Linux: Deployment [Guide](#)”. Especially these two settings:
 - nofile: 48000
 - nproc: 65536
 - **ETHERNET Device Queues:** Ensure that the ethernet device queues are greater than 0 for the CAS controller. We recommend multiple rx and tx queues for each ethernet adapter (current internal SAS Viya systems settings are 6 for both rx and tx). This applies to VM adapters as well. This setting helps prevent a line of micro-latencies from building up on ethernet queues. We determined there were detrimental performance effects caused by tx and rx queueing when the queues use a default value of 0.

In addition to all of the above information, consult the SAS Technical Paper “Checklist of SAS® Viya Administration Tasks” for information about what you need to consider prior to setting up a SAS Viya infrastructure.

GUIDANCE ON VIRTUALIZATION FOR SAS SERVERS

Now that we have discussed how to configure hardware for SAS®9 and SAS Viya applications, we need to address the particulars of implementing a virtual system (VMware ESX, Microsoft HyperV, Red Hat Enterprise Virtualization, and so on). In general, we need to address moving from a large physical system to several smaller virtual systems. The best practice for achieving optimal performance is to apply the same guidelines used for physical

system tuning to the virtual host operating system and the underlying hardware and I/O subsystem. Here are some additional tuning tips to follow:

- **Do not over-provision the underlying hardware in the physical frame.** This means that you should not overcommit your physical resources by allowing the virtual CPUs to grow beyond the physical backing, memory, or I/O connections that have been thickly and physically provisioned.
- **Ensure that virtual CPUs are backed with dedicated physical cores.** Ideally, you should have a 1:1 ratio between the virtual cores and the physical cores. Also, don't allow your virtual hypervisor to move the cores you are using randomly within the physical system as this can cause latency because the hypervisor will move only the processing application and not the data it is accessing in L1/L2 cache. This is because the hypervisor moves only the processing application and not the data it is accessing in the L1/L2 cache.
- **Make sure that the cores in the virtual CPU and L1/L2 cache are all on the same socket when creating the definition of a virtual CPU.**
- **If possible, connect the back-end storage with directly attached adapters to ensure optimal I/O throughput.** Any time you introduce conditions such as network traffic, additional IP traversals, and so on, you introduce the risk of encountering problems and performance degradations.
- **Consider the other virtual machines running on the same physical host and their impact on SAS applications.** For example, if you have a database running on another virtual host that consumes large volumes of I/O, this might impact the available I/O bandwidth inside the physical host that is being shared with SAS.

VIRTUALIZED STORAGE SOLUTIONS

In addition to examining the virtual CPU virtualization used for SAS sessions and servers, we can look at virtualization for storage arrays. Storage providers are offering a wave of new and advanced storage subsystems. These offerings promise virtualized, thin-provisioned, tiered, and intelligent storage that is easy to manage and that can reduce costs. For many random and mixed workload applications, these offerings deliver on their promises and perform well. Unfortunately, the SAS I/O workload profiles tend to violate some of the primary design assumptions underlying the configuration of these new systems. The SAS Global Forum paper “SAS® and the New Virtual Storage Systems” (Brown and Crevar 2013) addresses to the factors that must be considered when configuring these new storage systems.

Additional information about virtualization with SAS can be found in the SAS Global Forum paper “Virtualization: What does it mean for SAS®?” (Fisher and Thacher 2009).

GUIDANCE ON MONITORING HARDWARE

One thing we often see is that the use of SAS Viya will increase at your site over time. This growth can be caused by increases in the number of concurrent sessions running on the systems, in the size of the data being analyzed and reported, and in the types of SAS applications. Please note that the guidelines for setting up your hardware to run SAS are only guidelines, rather than detailed instructions. We recommend that you monitor your hardware on a regular basis to ensure that you do not run out of a computer resource and therefore cause SAS sessions to perform poorly or terminate unexpectedly.

There are several papers being presented at SAS Global Forum 2020 that discuss ways to monitor your SAS Viya infrastructure. Two that might be of interest are “Diagnosing the Most Common SAS Viya Performance Problems” (Kuell 2020) and “Solving Performance Problems in a SAS® Viya® Multi-Tier Deployment” (Blake 2020).

CONCLUSION

There are many things that are needed to maintain happy SAS Viya users, many of which have been discussed in this paper. However, one that we have not discussed high availability (HA), failover, and disaster recovery. SAS does recovery policy for SAS Viya (see References). We mention briefly that supporting HA by spitting CAS environment across two data centers is not ideal if the network between the two data centers is not robust enough to handle the movement of data and communications between the SAS Viya nodes.

We hope this paper has helped you understand the compute resources that are needed for SAS Viya, along with how to best configure and monitor these resources. We will update this paper each year as we learn more about SAS Viya usage from our customers and users.

REFERENCES

- "Analytics Life Cycle." SAS Institute Inc. Available <https://smportal.sas.com/software/analytics-life-cycle/Pages/default.aspx#quickstart>. Accessed February 21 2020.
- Schneider, Mark. 2019. "Comparing SAS® Viya and SAS® 9.4: How their Features Complement Each Other". *Proceedings of the SAS Global 2019 Conference*, Cary, NC: SAS Institute Inc. Available <https://www.sas.com/content/dam/SAS/support/en/sas-global-forum-proceedings/2019/3503-2019.pdf>
- Brown, Tony. 2019. "Engineering CAS Performance Hardware Network, and Storage Considerations for CAS Servers" *Proceedings of the SAS Global Forum 2019 Conference*. Cary, NC: SAS Institute Inc. Available <https://www.sas.com/content/dam/SAS/support/en/sas-global-forum-proceedings/2019/3351-2019.pdf>
- [Brown, Tony, and Margaret Crevar. 2016. "Best Practices for Configuring Your I/O Subsystem for SAS9 Applications." *Proceedings of the SAS Global Forum 2016*. Cary, NC: SAS Institute Inc. Available <http://support.sas.com/resources/papers/proceedings16/SAS6761-2016.pdf>](http://support.sas.com/resources/papers/proceedings16/SAS6761-2016.pdf)
- Collum, Rob. 2019. "Seriously Serial or Perfectly Parallel Data Transfer with SAS® Viya®." *Proceedings of the SAS Global Forum 2019 Conference*. Available <https://www.sas.com/content/dam/SAS/support/en/sas-global-forum-proceedings/2019/3479-2019.pdf>
- [Marson, Barry. "Optimizing SAS on Red Hat Enterprise Linux \(RHEL\) 6 & 7" Raleigh, NC; Red Hat Inc. Available \[https://access.redhat.com/sites/default/files/attachments/optimizingsasonrhel6and7_2.pdf\]\(https://access.redhat.com/sites/default/files/attachments/optimizingsasonrhel6and7_2.pdf\). Accessed February 21 2020.](https://access.redhat.com/sites/default/files/attachments/optimizingsasonrhel6and7_2.pdf)
- Stern, David. SAS Institute Inc. "Checklist of SAS® Viya Administration Tasks." Cary, NC: SAS Institute Inc. Available <https://support.sas.com/resources/papers/checklist-sas-viya-administration-tasks.pdf>. Accessed February 21, 2020.
- Brown, Tony, and Margaret Crevar. 2013. "SAS® and the New Virtual Storage Systems." *Proceedings of the SAS Global Forum 2013 Conference*. Cary, NC: SAS Institute Inc. Available <http://support.sas.com/resources/papers/proceedings13/487-2013.pdf>
- Fisher, Karl, and Clarke Thacher. 2009. "Virtualization: What does it mean for SAS®?" *Proceedings of the SAS Global Forum 2009 Conference*. Cary, NC: SAS Institute Inc. Available <http://support.sas.com/resources/papers/sgf09/347-2009.pdf>

SAS Institute Inc. "SAS Disaster Recovery Policy for SAS® Viya® 3.4." Available <https://support.sas.com/en/technical-support/services-policies/disaster-recovery-policy-viya.html>. Accessed February 21, 2020.

RobCollum. 2019. "4 Rules to Understand CAS Management of In-Memory Data" *SAS Communities – Administration*. <https://communities.sas.com/t5/SAS-Communities-Library/4-Rules-to-Understand-CAS-Management-of-In-Memory-Data/ta-p/588652>

RobCollum. 2019. "Provisioning CAS_DISK_CACHE for SAS Viya" *SAS Communities – Administration*. <https://communities.sas.com/t5/SAS-Communities-Library/Provisioning-CAS-DISK-CACHE-for-SAS-Viya/ta-p/603689>

RobCollum. 2020. "Understanding Network Binding Variables in SAS Viya 3.5" *SAS Communities – Administration*. <https://communities.sas.com/t5/SAS-Communities-Library/Understanding-Network-Binding-Variables-in-SAS-Viya-3-5/ta-p/627000>

ACKNOWLEDGMENTS

Many thanks to everyone who helped review this paper. There are too many to list.

RECOMMENDED READING

- [SAS Help Center Downloadable PDFs for SAS Viya 3.5 Administration](#).

This web site enables you to download a current snapshot of the SAS Viya administration documentation. You should check back often, because these guides are periodically updated.

CONTACT INFORMATION

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

Margaret Crevar
SAS Institute Inc
Margaret.Crevar@sas.com
www.sas.com

SAS and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration.

Other brand and product names are trademarks of their respective companies.