Paper SAS4296-2020

Diagnosing the Most Common SAS® Viya® Performance Problems

Jim Kuell, SAS Institute Inc.

ABSTRACT

Applications and the scalable computing environments in which they run have grown in complexity with more advanced technologies. With the mixture of virtual machine, cloud, and emerging container environments, diagnosing the causes of performance issues can be difficult. Relying on significant experience from the SAS® Performance Lab, this paper presents the most common SAS® Viya® performance problems and methods for diagnosing and correcting them.

INTRODUCTION

SAS® Viya® is an extremely performant analytics engine that is designed to provide quick and accurate analytical insights in even the most complex environments. However, as environments grow in complexity, so do the number of factors that need to be taken into consideration when architecting for performance. When performance issues arise, properly diagnosing and correcting them can be a very tedious process, especially if you don't know what to look for.

This paper provides an overview of the SAS Performance Lab's (SPL) methodology for diagnosing SAS Viya performance problems and discusses the most common causes of these problems at customer sites. The hope of this paper is that the information shared will help lead to the expedited resolution of performance issues that arise.

OUR METHODOLOGY

While the forensic information we gather for SAS Viya is relatively similar to what we gather for SAS®9, the process of analyzing that information is very different. Between all the microservices, application layers, physical and virtual infrastructures, file systems, and so on, there are *many* more moving parts and areas to analyze with SAS Viya.

For instance, CAS procedures and actions are now multi-threaded and running in parallel, but there is a cost for launching these threads, especially across CAS workers. There is a point that you can cross where you are over-threading a task-set and actually hindering performance. Additionally, over-distributing data (spreading too little data over too many systems) can also extremely encumber performance. As a general rule of thumb, we typically recommend data that is distributed should be at least 1 GB in size for each host.

The items listed in this section need to be carefully gathered so that they can be overlaid and analyzed to get to the bottom of an issue. Correctly gathering the information is the most important step of the SAS Performance Lab's diagnosis process. The analysis process is very fluid and requesting more information/monitoring during this phase is very common. Active cooperation from the customer is crucial to quickly and effectively resolve performance issues.

PROBLEM DEFINITION

The first piece of information needed is a very detailed definition of the performance problem. Finding the origin of a problem is impossible without first clearly understanding what the problem is.

This definition needs to not only include the problem (the user's or admin's perception of what's happening), but what is actually occurring on the system(s). For example, is there a single table that's being operated on or does each user have a separate table that they're loading and trying to process? Is the job running slower than previously? Does it hang up and not respond?

APPLICATION DEFINITION

Different applications have varying impacts on the performance of a system. It is important to compile and map out a list of all applications (SAS Viya and others) that interact with each of the problematic compute systems. This includes the interface that is used to run the SAS jobs in question (that is, SAS® Enterprise Guide®, SAS® Studio, batch, and so on).

In this list, you should identify the specific SAS Viya applications (for example, SAS® Visual Analytics, SAS® Visual Statistics, and so on) that appear to be exhibiting the performance problems.

INFRASTRUCTURE DEFINITION

Many performance problems can be traced back to issues within the infrastructure, either hardware- or software-based. This includes the servers, networks, operating system-level tunings, and I/O subsystems. Gathering and understanding this information might require the assistance of the Systems, Network, or Storage Administrators. Collecting the following information for each of the problematic systems helps us begin creating a detailed infrastructure mapping:

1. Server and Network Information
    a. Manufacturer and model of the systems
    b. Virtualization or containerization software being used, if any
    c. Model and speed of the CPUs
    d. Number of physical CPU cores
    e. Amount of physical RAM
    f. Network connection type and speed
2. OS-Level Information
    a. Operating system version
    b. File systems being used for both permanent SAS data files (SAS Data) and temporary SAS data files (CAS Disk Cache and SAS Work)
    c. Source data locations (for example, SAS data files, external database, Hadoop, and so on)
    d. OS tuning parameters and settings
3. I/O Subsystem Information
    a. Manufacturer and model number of the storage array and/or devices
    b. Storage types and physical disk sizes, as well as any relevant striping information (that is, RAID, and so on)

c. Types of connections used (for example, NICs, HBAs, and so on), the number of cards and ports, and the bandwidth capabilities of each (for example, 8 Gbit, 10 Gbit, and so on)
d. I/O-related tuning parameters and settings

The SPL has also developed a tool that gathers and packages the output of various operating system-level commands and files. This tool is called the *RHEL Gather Information Script* and a link to download this tool can be found in the *Output From Tools* section below. The output of this tool provides a more detailed look at the OS-level information and is used as a supplementary knowledge base to the information that is listed above. It gives us a closer look at specific OS tunings, user `ulimit` information, logical volume configurations, and much more.

SAS LOGS

The SAS log is typically where we begin our investigation.

SAS logs from the jobs that are surfacing the performance issues are vital to the diagnosis process. They contain performance metrics and session options from a SAS job run that help us narrow down the source of an issue. If the SAS job in question ever ran without performance issues, it would be extremely beneficial to also collect the log from that run.

We also look at the logs for any applications that are exhibiting performance issues. Turning on debug for these logs is extremely beneficial. SAS Technical Support can provide specific instructions on how to turn on debug for each application. They can also provide several additional SAS log settings for individual jobs that print out helpful metrics and configurations (for example, the 'metrics=true' CAS session option, and so on).

We cannot determine what the bottleneck is from the SAS logs alone. We need to corroborate this information with the other hardware monitors and system information we collect. Since the log tells us the exact time that the step was executed, we can overlay that with output from the other tools and isolate the data from that specific time frame only. We can then use the step's metrics from the SAS log to give us a better idea of where to look in the output.

In addition, the SAS log also contains information about which file systems are used by different caslibs and data stores. We can then examine the file systems that are used by problematic steps and actions to determine whether they are configured correctly.

Because of the large amount of moving parts with SAS Viya, analyzing logs has proven to be much more complex than it was for SAS®9. Because of this, we ask customers with performance issues to collect the logs and submit them to their SAS Technical Support track. We then work directly with SAS Technical Support to analyze them to get to the bottom of the issue(s) at hand.

OUTPUT FROM TOOLS

There are a number of diagnostic tools that SAS Technical Support will direct you to employ, with explicit directions and help, when necessary. The output files created by the tools in this section provide a vast amount of information about many different aspects of the environment in question. Further details about what each tool collects can be found in the links provided below.

Note that some of the tools listed below may not have the same functionality when used in containerized environments. We have other tools and methods of diagnosis that we can work with you to employ if you are experiencing performance issues in a containerized environment.

- RHEL Gather Information Script [1] – This tool was developed by SAS and gathers and packages the output of various OS-level commands and files. It provides the information needed to validate that the operating system is correctly tuned as well as other useful system configurations. It should be run on all problematic systems. For more information, see http://support.sas.com/kb/57/825.html.

- SAS Viya Perf Tool [1] – This is a free, standalone tool created by SAS. It runs a series of tests on all of the SAS Viya nodes, including testing the available throughput of the network between the nodes as well as the file systems used by each node for both permanent SAS data files and CAS Disk Cache. There are several available options in the tool's configuration file that allow you to choose between several types of tests and test options. An announcement will be made in the Administration SAS Communities site (https://communities.sas.com/t5/Administration-and-Deployment/bd-p/sas_admin) when this tool is made available. More information and a link to download the tool will be located here: http://support.sas.com/kb/53/877.html.
- IBM **nmon** script [1] – This free tool is our preferred hardware monitor for RHEL systems. It collects a large amount of information from the system kernel monitors, and its output can later be converted into a graphical Microsoft Excel spreadsheet using the `nmon_analyzer` tool. The `nmon_analyzer` splits the information into a series of tabs and makes it much easier to read. Nmon should be run during the SAS jobs in question. Overlaying it with both the SAS log and the other tools mentioned here allows us to much more easily hone in on what is causing the performance issue at hand. For more information, see http://support.sas.com/kb/48/290.html.
- Gridmon [2] – This is an administration and monitoring utility that is included in SAS Viya deployments. It not only allows you to stream and display data from CAS server nodes in real time, but it allows you to record this information to a file. This file can then be sent to SAS Technical Support for review. The information captured by `gridmon` recordings often plays a vital role in getting to the bottom of SAS Viya performance problems. For more details about how use `gridmon` and what information is captured in its recordings, see https://go.documentation.sas.com/?docsetId=calserverscas&docsetTarget=n03061viyaservers000000admin.htm&docsetVersion=3.5&locale=en.
- Tkgridperf [2] – This tool is used to isolate network issues between nodes in SAS Viya deployments. It utilizes several methods of sending chunks of data between nodes and the results can be used to evaluate if the entire cluster is working well together. Data is sent from the CAS Controller and is cascaded down to the CAS Workers in a tree/branching like pattern, with the branching pattern repeated through the CAS Workers. This tool should be run periodically over time and compared to each other to assess the status of the network and whether or not there has been any degradation. Running it a single time is generally not as helpful or easy to assess as periodic runs. Note that several features have been added in SAS Viya 3.5 that are not available in earlier versions.
- sas-peek [2] – This tool is a metric collection utility that's included with SAS Viya deployments. It finds and queries various components of SAS Viya and produces JSON output containing the metric data. The metric data it produces includes metrics for the host system (such as CPU, memory, network, filesystems, and I/O), processes, CAS servers, SAS microservices, RabbitMQ, and Postgres.

  Because the sas-peek command runs on every machine in a deployment, by default only local resources are reported to avoid duplication of metrics. Multiple levels of detail are supported for each type of metric collected. Although the sas-peek command is normally executed by the sas-ops-agent command, it can be run manually. For more information, see https://www.sas.com/content/dam/SAS/support/en/sas-global-forum-proceedings/2020/4214-2020.pdf.

1 – Free, standalone tools that do not require a SAS® installation or license.

2 – Tools that are shipped with and included in SAS Viya installations.

When correlated with the information listed in the previous sections, the output from these tools is often enough to track down the root cause of most performance issues. Additional tools are available and used on a case-by-case basis when more complex issues arise.

THE MOST COMMON CAUSES OF PERFORMANCE PROBLEMS

SAS Viya has proven to vastly reduce the run time of many data manipulation and statistical procedures through parallelization. However, this requires a consistently well-performing infrastructure. The SAS infrastructure and the many factors it's comprised of is the most common cause of performance problems with SAS Viya.

Before implementing or changing a SAS infrastructure, you must fully understand every aspect of the environment and its expectations: specific SAS applications, number of users, data sizes, anticipated growth, security requirements, and so on. You must also have a good understanding of the SAS workload requirements and the hardware infrastructure needed to meet the service-level agreements (SLAs).

You may be asking yourself, "How do I know if my systems are not performing well?" SAS Technical support tracks for performance are typically opened when the customer is experiencing one of the following issues: a poor user experience; jobs or sessions are taking too long to run; jobs or sessions fall over and never finish; jobs or sessions "freeze-up" and become unresponsive. We begin our investigation when one of these scenarios occurs.

While every situation varies in complexity, the following causes account for the most common SAS Viya performance problems.

NETWORK

Insufficient network bandwidth for intercommunication between SAS Viya nodes is the most common performance problem that we see. The SAS Viya systems – CAS Controller, CAS workers, SAS Programming Run-Time node, CAS Microservices node, RabbitMQ/Postgres node – are all very chatty and there is constant and heavy communication between them. In addition, data is often transferred from node to node for parallel operations or backing store availability (CAS Disk Cache). Because of this, we strongly recommend a minimum of a 10 Gigabit NIC for inter-node communication. Higher bandwidth NICs, LAN switches, and fabric may be required depending on the usage patterns and data sizes.

PHYSICAL LOCATION

It is very important that every component of a SAS infrastructure is co-located in the same physical location (that is, close proximity in network segments, same cloud availability zones, and so on). This includes the source data files, storage, systems, authentication, and so on. These components should always be located on the same subnet. Serious performance degradation occurs when they are not and additional networks (specifically WAN) come into play. For more information, see https://communities.sas.com/t5/Architecture/Does-It-Matter-Where-the-Various-Components-of-Your-SAS/m-p/483426.

BANDWIDTH TO EXTERNAL STORAGE

SAS Viya has heavy I/O throughput needs, as does SAS® 9, and this is often not accounted for when infrastructures are architected. You need to make sure you have sufficient bandwidth between the data and the servers. Our minimum recommended I/O throughput is 100 MB/sec/physical core for file systems that house permanent data and 150 MB/sec/physical core for temporary work-space file systems.

It is also important to note that your I/O infrastructure is only as fast as its slowest component. You should consider the peak I/O throughput requirements of the environment and work with your storage vendors to ensure that each segment of your I/O infrastructure can meet those demands.

For network-attached storage, you may want the NIC that is used to transfer to/from external storage to be separate from the NIC used for inter-node communication between the SAS servers. The amount of data loaded into CAS is often very large, so you may need multiple bonded network adapters in each node to achieve the throughput that is required. How much throughput is needed greatly varies from site to site. This depends on the service-level agreements (SLAs) of the system's jobs and is directly driven by the amount of time you can afford to spend loading the input data for CAS actions from disk into memory.

MEMORY

SAS Viya is an in-memory processing system that utilizes memory-mapped files, and in-memory processes need to stay within virtual memory limits to perform optimally. Because of this, the CAS Controller and Workers need an ample amount of memory for SAS Viya to be performant. Correctly sizing memory capacity is extremely important so that CAS can run at memory speeds as much as possible.

The amount of memory needed varies greatly from site to site, depending on your data sizes, applications, workload characteristics, and so on. We recommend that the size of memory is a minimum of 2x the size of the total concurrent incoming data. We've found that this is typically a good starting place for memory capacity planning, and it can be expanded from here. Tasks that include significant amounts of data manipulation, for example, may require a memory size of 3x the size of concurrent incoming data. It's very important that you perform a detailed workload analysis before estimating your required memory capacity.

CAS DISK CACHE

There are many important factors that need to be taken under consideration when designing your CAS Disk Cache file system. For more information on how to properly architect this file system, please review the papers listed in the references section below.

CAS Disk Cache serves as an on-device extension of memory-maps, providing extra space for in-memory data, as well as a backing store for blocks of data that need to be replicated.

Although CAS utilizes the incredible efficiency of the CAS Disk Cache file system, performance takes a big hit when you do not have enough RAM and exceed memory space. This is because the extremely fast processing speed of memory is slowed to the maximum throughput rate of the CAS Disk Cache file system and the performance of the devices it resides on. It's important to note that CAS Disk Cache is still utilized even if you have enough memory to process your data. How much it's used depends on the type of data being used and what exactly is being done with it.

The default location of CAS Disk Cache is /tmp. We highly recommend changing the location of this for several reasons: 1) The danger of filling up /tmp – if this occurs, the operating system cannot run correctly and everything on the node stops; 2) /tmp is typically not a very performant file system. The ideal scenario would be to change the location of CAS Disk Cache to local drives that are designed for performance. The minimum throughput that we recommend for this file system is 150 MB/sec/physical core.

We find that sizing this file system isn't always a straightforward task. We typically recommend starting with 1.5x – 2x the size of RAM. However, it's important that you are ready to grow its size if you start executing your workloads and find that more space is needed. In many situations we find that 3x the size of RAM or more is required for CAS Disk Cache, depending on the specific actions and data sizes. Your capacity requirements may vary considerably.

VIRTUALIZATION

Virtualization is a very common practice in SAS shops around the world. Virtualized environments can function very well with minimal overhead if they are configured and provisioned correctly. SAS does not work well when it's located on a thinly provisioned, virtualized server farm with shared everything (cores, memory, and storage arrays), so it is important to provide it with dedicated hardware resources.

Many performance problems can also be traced back to virtualized environments with an underlying hardware infrastructure that cannot meet the I/O throughput required by the SAS. The underlying hardware needs to be set up and tuned like it always has been with SAS® 9, and this is often not the case.

Another common issue with virtualized environments is host definition placement of CPU socket arrangements. Virtual machines can encounter non-uniform memory access (NUMA), which can degrade performance, if the virtual CPUs (vCPUs) span multiple physical sockets. A good practice is to keep all the cores in your virtual system on a single socket within the underlying hardware.

The specific application or applications, data sizes, and workload characteristics must be carefully analyzed to determine the appropriate sizing and configuration of a virtualized infrastructure.

CONCLUSION

When performance issues arise, knowing what to look for and where to begin can be extremely difficult. The purpose of this paper is to assist in that process by providing an overview of the forensic process that the SAS Performance Lab uses when diagnosing performance problems, as well as the most common causes, what to look for, and methods for correcting them. This information can be used by customers to either begin a path toward self-diagnosis, or preemptively gather all the required information before opening a track with SAS Technical Support, leading to an expedited resolution.

REFERENCES

SAS Institute Inc. "SAS Note 42197: A list of papers useful for troubleshooting system performance problems." http://support.sas.com/kb/42/197.html.

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. https://www.sas.com/content/dam/SAS/support/en/sas-global-forum-proceedings/2019/3351-2019.pdf.

Crevar, Margaret. 2020. "Important Performance Considerations When Moving SAS® to a Public Cloud." *Proceedings of the SAS Global Forum 2020 Conference*. Cary, NC: SAS Institute Inc. https://www.sas.com/content/dam/SAS/support/en/sas-global-forum-proceedings/2020/4312-2020.pdf.

Ellington, Bryan. 2020. "SAS® Viya® Monitoring Using Open-Source Tools." *Proceedings of the SAS Global Forum 2020 Conference*. Cary, NC: SAS Institute Inc. https://www.sas.com/content/dam/SAS/support/en/sas-global-forum-proceedings/2020/4214-2020.pdf.

Venkataramani, Meera. 2020. "Essential Performance Tips for SAS® Visual Analytics." *Proceedings of the SAS Global Forum 2020 Conference*. Cary, NC: SAS Institute Inc. https://www.sas.com/content/dam/SAS/support/en/sas-global-forum-proceedings/2020/4408-2020.pdf.

ACKNOWLEDGMENTS

A special thank you to Margaret Crevar and Tony Brown for their contributions to this paper.

CONTACT INFORMATION

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

Jim Kuell
SAS Institute Inc.
Jim.Kuell@sas.com
http://www.sas.com