

Important Performance Considerations When Moving SAS® to a Public Cloud

Margaret Crevar, SAS Institute Inc., Cary, NC

ABSTRACT

When choosing a hardware infrastructure to run your SAS applications, you need a solid understanding of all the layers and components of the SAS infrastructure. You also need the ability not only to successfully run the software but to optimize its performance. Finally, you need an administrator to configure and manage the infrastructure. This paper talks about important performance considerations for SAS®9 (both SAS® Foundation and SAS® Grid Manager) and for SAS® Viya® when hosted in any of the available public clouds—Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform, to name a few. It also provides guidance on how to configure the cloud infrastructure to get the best performance with SAS.

Disclaimer: We strongly encourage you to take the advice in this paper and work with your local public cloud teams to make sure the instances you decide to use are available in the closest region and that you understand their costs. In addition, any advice in this paper is based on the information we have at the time of publishing this paper (March 2019).

INTRODUCTION

Many SAS customers are making the decision to move their current SAS applications from their on-premises data centers to a public cloud. The hype around public clouds portrays this as a very simple task that saves SAS customers a lot of money.

It should be noted that the information discussed in this paper is based on what is available from the public clouds and SAS' experience with the public clouds at the time of its writing. Public cloud offerings are constantly changing. Therefore, it is in your best interest to understand the rationale used in the selection process and to consider what was done as a point-in-time design.

However, there is a lot of planning needed and, depending on the requirements from the SAS customer, the price might not be cheaper than on-premises hosting. The chief reason is that if IO throughput is crucial to the success of your SAS applications in the public cloud, you might need to provision more cores and disk space capacity to ensure the success of SAS public cloud deployments. The reasons for this are explained in this paper.

BEFORE YOU START

As mentioned in the introduction, a good understanding of the SAS workload requirements, along with the hardware infrastructure required to meet the service objectives (SLAs), specifically the time to complete the task, is crucial. For existing SAS customers, the following questions help guide that examination:

- Are there SAS jobs that need to execute within a certain time frame? Are you expecting your SAS jobs to execute in the same time frame—or faster than—they are currently running in your existing data center? If so, a determination of the IO throughput required for each file system being used must be made. It must be determined if this same IO throughput can be achieved in the public cloud.
- Where is the source data for the SAS jobs located? Does this data already reside in the public cloud of choice? If not, the amount of time required to move data to the cloud space where SAS is executing must be determined. This added time will affect the SLA of the jobs that consume off-cloud data.
- Is the customer's IT staff willing to do stand-up authentication in the public cloud?

- What security is needed for the data and/or SAS code?

The answers to these questions and fact-finding need to be fully understood so that the correct hardware and storage are selected from the available public cloud offerings. There are many different hardware and storage types. Some are hardware equipped for the heavy analytical and large sequential IO that SAS 9 does. Others are better equipped for the in-memory needs of SAS Viya. It is important to understand the workload profile of the customer's SAS application(s) to ensure that correct hardware and storage selections (cloud server and storage types) are made for the best performance. Please note that to get the best achievable performance, the least expensive hardware and storage types from the public cloud offerings might not be suitable. For example, the customer might require server and storage instances with more physical cores than required for computing needs and/or more storage capacity than the initial sizes needed to acquire the maximum IO bandwidth available for their SAS application(s).

Now let's talk about what needs to be considered to ensure that you can configure the hardware infrastructure in the public cloud to perform as optimally as possible. These things include the following:

- what server instance type to use for your SAS 9.4, SAS Viya, or hybrid SAS infrastructure
- what storage type (for both persistent and nonpersistent storage) to use
- if you are deploying SAS® Grid Manager, what shared file system to use
- where to place temporary (SASWORK/UTILLOC and CAS_Disk_Cache) and permanent (SASDATA) data to be used by SAS
- where to place the SAS clients that will be used
- where to place authentication tools
- whether high availability and security are required

WHAT INSTANCE TYPE TO USE

In SAS 9 and SAS Viya infrastructures, there are several SAS server types and uses. Each has different and specific requirements for CPU, IO throughput, and memory provisioning. We will list each SAS server type and discuss its provisioning requirements. Please remember that most public cloud instances list CPUs as virtual CPU(s). These CPUs might actually be hyperthreaded (two threads per CPU core). You need to understand if the vCPU includes hyperthreads so that you can ensure you have the correct number of physical cores for SAS.

You might have to use an instance with more physical cores in it than your workload requires. This is because a higher CPU count machine might be required to obtain a dedicated network interface card (NIC) of sufficient bandwidth to maximize IO to and from off-server cloud storage. Server instance types come in set models, with set CPU counts, set NIC installation, RAM, and so on. To get a dedicated NIC or the IO bandwidth required for your SAS workload through the NIC to the off-server storage, you might have to upgrade to a larger server type (with more CPU and RAM than needed for computing). When calculating NIC IO bandwidth capacity to drive storage, be aware that sharing a NIC in a cloud server host with other multi-tenant applications residing on the same physical server might result in inferior performance that occurs randomly. This is especially true for virtualized cloud host instances.

When setting up the instances, please make sure that all the instances and components of your SAS infrastructure (fat clients, data, authentication software, SAS products, and so on) are in the same cloud space (region, area zone, placement group, and so on). Failure to do so results in an additional WAN connection that severely impacts performance.

SAS 9 SYSTEMS

These are general guidelines. We suggest that you have a detailed workload assessment done to determine which hardware is needed to support the usage of SAS 9 in the public cloud.

SAS Compute Tiers Including SAS Grid Node

These servers need fast CPUs for processing data, a minimum of 8 GB of physical RAM per physical core, and robust IO throughput (especially to SAS WORK and SAS UTILLOC):

- Amazon
 - I3 series - The primary reason is the high internal IO bandwidth from striped NVMe SSD drives for SAS WORK and UTILLOC file systems.
 - C5N series – This series has a larger NIC card, but does not have the amount of RAM per physical core that is recommended
- MS Azure
 - GS5 series
 - ESv3 series
 - M series – if you would like the largest amount of RAM per physical core; however, these are not available in every MS Azure region.
 - DSv3 series – these have the smallest ephemeral disks.
- Google
 - N1-standard series

Shared File System Storage Required for SAS Grid

These servers need robust IO throughput to the permanent storage that will support shared file systems. These instances will also need a minimum of 8 GB of RAM per physical core:

- Amazon
 - R5 or R5d series
- MS Azure
 - GS5 series
 - ESv3 series
- Google
 - N1-standard series

SAS Mid-Tier and Metadata Servers

These servers do not require computing-intensive resources or robust IO bandwidth, but they do require access to more memory than the SAS computing tiers. The recommendation is a minimum of 24 GB of physical RAM or 8 GB of physical RAM per physical core—whichever is larger:

- Amazon
 - R5 or R5d series
 - M series
 - X1e series
- MS Azure
 - M series - if you would like the largest amount of RAM per physical core; however, these are not available in every MS Azure region.
 - ESv3 series
 - DSv3 series – these have the smallest ephemeral disks.
- Google
 - N1-highmem series

SAS VIYA SERVERS

Following is a list of several machines to use for the most robust SAS Viya offering. Some of the SAS Viya offerings (single, non-distributed node going against small data from a source other than Hadoop) will not need all the servers listed here. These are general guidelines. We suggest that you have a detailed workload assessment done to determine which hardware is needed to support the usage of SAS Viya in the public cloud.

CAS Nodes – Minimum of Three

These servers require fast CPUs for processing data, enough physical RAM to hold all the data files to be analyzed by all the concurrent SAS Viya users, and robust IO throughput (especially to CAS_Disk_Cache). If you are not sure how much data will be accessed at any given time, but you know your SAS users will be accessing files in the 100s of gigabytes in size, we recommend 64 GB of RAM per physical core (more if you have large files that will reside in memory):

- Amazon
 - I3 series – if your SAS Viya usage will need fast access to CAS_Disk_Cache
 - X1e series – if you would like the largest amount of RAM per physical core
 - R5, R5d, and R5a series
- MS Azure
 - M series – if you would like the largest amount of RAM per physical core; however, these are not available in every MS Azure region.
 - GS series
 - ESv3 series
 - DSv3 series – these have the smallest ephemeral disks.
- GS series Google
 - N1-highmem series

MicroServices Node

These servers do not require high computational speed or power. To run all SAS® Visual Analytics products with SAS Viya, you will need at least 96 GB of RAM in your MicroServices Node:

- Amazon
 - R5/R5d/R5a or X1e series
- MS Azure
 - M series
- Google
 - N1-highmem series

SAS® Programming Run-Time Node

This server will run SAS 9 code based on the application. This server needs fast CPUs for processing data (at least 16 GB of RAM per physical core) and robust IO throughput (especially to SAS WORK and SAS UTILLOC):

- Amazon
 - I3 series – the primary reason is the high internal IO bandwidth from striped NVMe SSD drives for SAS WORK and UTILLOC file systems.
- MS Azure
 - GS series
 - ESv3 series
 - DSv3 series – these have the smallest ephemeral disks.
- Google
 - N1-standard series

WHICH STORAGE TYPE TO USE

The public clouds have multiple types of permanent and temporary data storage offerings for use with SAS. Discussed here are findings from general field experience and lab testing with SAS in public cloud spaces.

Permanent SAS Data Storage Will Persist Through a Reboot/Restart

Permanent SAS data storage is used for SAS 9.4 SAS data files and SAS Viya CAS tables:

- Amazon

- Elastic Block Storage (EBS) – Our recommendation is to stripe together a minimum of 4 EBS volumes for IO bandwidth aggregation. Please see the [EBS Storage guide](#) on the Amazon website for more details.
 - EBS ST1 (Throughput Optimized HDD) Storage – designed for large block sequential IO. A 12.5 TB volume can sustain 500 MB/second (This is a published value from Amazon). If your volume size is less than this, you will only get 500 MB/second total bandwidth during your burst window.
 - EBS IO1 (Provisioned IOPS SSD) storage can also be used.
 - Other EBS storage types - like general-purpose SSD (gp2), and cold storage (sc1) - should not be used for permanent SAS 9 data files.
- S3 storage – please review the SAS Usage Note 63001 (<http://support.sas.com/kb/63/001.html>) that discusses ways to get additional functional and better performance to S3 with SAS.
- MS Azure
 - Premium Storage Disk Type. You will need to review the IO throughput per storage disk type (<https://docs.microsoft.com/en-us/azure/virtual-machines/windows/premium-storage-performance#premium-storage-disk-sizes>) to determine how many disks to stripe together to obtain the minimum of 100 MB/physical core/second. To get the minimum IO throughput for a 16 physical core instance, you will need to stripe multiple Premium Storage disks eleven P20s, eight P30s, seven P40s/P50s, four P60s, or three P70s/P80s together.
- Google
 - Use the storage with the highest IO throughput. You should stripe at least 4 (preferably 8) together for each instance.

Temporary SAS Data Storage Will Not Persist Through a Reboot/Restart

Temporary storage is most commonly used for SAS WORK, SAS UTILLOC, and CAS_Disk_Cache because this data does not need to persist through reboots and restarts. When placed on striped NVMeS, SAS WORK and SAS UTILLOC can share a single file system:

- Amazon
 - Internal SSD devices striped together with RAID0. We would prefer you use instances with NVMe (ephemeral) devices with high bandwidth, low latency, and sequential IO - all of which are ideal for temporary SAS data. *
- MS Azure
 - Internal SSD devices, at the present time, *cannot* be striped together with RAID0. However, you can use a single SSD device for SAS WORK. *
- Google
 - Internal SSD devices striped together with RAID0 *

* Note: If there is not enough capacity using internal drives, please follow the advice for permanent SAS data files listed earlier.

IF DEPLOYING SAS GRID MANAGER, WHICH SHARED FILE SYSTEM TO USE

SAS Grid Manager requires a shared file system for the permanent files being shared by all the SAS Grid compute nodes. Here are several shared file systems that have been tested with SAS Grid Manager in several public cloud infrastructures:

- DDN Lustre - For optimal performance, you will need a *minimum* of four IO robust instances for your Lustre shared file system. A paper from April 2015 that discusses how we configured Lustre to work in AWS can be found here: <https://support.sas.com/rnd/scalability/grid/SGMonAWS.pdf>.

Please note that you must mount your Lustre file systems with this option (*flock*) in order for Lustre to work optimally with SAS, especially if you plan to host your SAS binaries on Lustre.

Also note that you need to understand that you are unable to expand the size of your Lustre file system. You will need to create a larger Lustre file system and copy data from the old system to the new one.

You can use your own copy of Lustre, or you can use a copy from the public cloud marketplace. Links to the marketplace offerings are listed here:

- Amazon –
 - Premier support: https://aws.amazon.com/marketplace/pp/B07FTX9GBV?qid=1548776061090&sr=0-1&ref=srh_res_product_title
 - Self support: <https://aws.amazon.com/marketplace/pp/B07FTSL9VX>
 - Microsoft Azure – not available at the writing of this paper
 - Goggle – not available at the writing of this paper
- IBM Spectrum Scale (formerly known as GPFS) – For optimal performance, you will need a *minimum* of four IO robust instances for your Spectrum Scale shared file system. A paper from April 2018 that discusses how to configure Spectrum Scale in AWS can be found here: <https://support.sas.com/resources/papers/performance-tuning-considerations-amazon-web-services-sas-9-4-ibm-spectrum-scale.pdf>.

You can use your own copy of Spectrum Scale, or you can use a copy for the public cloud marketplace. Links to marketplace offerings are listed here:

- Amazon - https://aws.amazon.com/marketplace/search/?filters=vendor_id&vendor_id=2f01b596-7972-4714-be52-b3709033137e
 - Microsoft Azure - not available at the writing of this paper
 - Goggle – not available at the writing of this paper
- Amazon FSx for Lustre – This is a new shared file system only available on AWS. Information about this file system can be found here: <https://aws.amazon.com/fsx/lustre/>.

Please note that AWS sets up the Lustre file systems with these mount options (*rw,seclabel,lazystatfs*). As you can see from the information here, these are not the mount options that SAS prefers, so you will need to unmount the FSx for Lustre file systems that were set up for you by Amazon and remount them with the *flock* parameter.

Also note that you need to understand that you are unable to expand the size of your Lustre file system. You will need to create a larger Lustre file system and copy data from the old system to the new one.

- Amazon Elastic File System (EFS) – This file system is only available in AWS. Please note that the maximum IO throughput is currently 250 MB per second per instance, refer to Amazon EFS limits (<https://docs.aws.amazon.com/efs/latest/ug/limits.html#limits-efs-resources-per-account-per-region>) for the latest information on limits. You will need to have multiple EFS file systems per instance to overcome this IO throughput limitation, but you also need to understand that there is a limitation on the single NIC in each AWS EC2 instance. These file systems *cannot* be striped together.

Please note that AWS recently increased the number of hard file locks from 87 to 512 for these shared file systems.

- [Azure NetApp](http://support.sas.com/resources/papers/proceedings16/SAS6761-2016.pdf) – This NFS-based file system is only available in Azure. Please note that all the concerns with heavy WRITES to an NFS file system apply here. For more details about NFS and SAS 9.4 can be found starting on page 4 in this paper: <http://support.sas.com/resources/papers/proceedings16/SAS6761-2016.pdf>.

WHERE TO PLACE VARIOUS SAS COMPONENTS

For optimal performance, it is very important to place all of the components of your SAS infrastructure in the same public cloud region, placement group, and availability zone. This includes the clients, SAS binaries, SAS tiers, source data, authentication tools, and so on. If you do not place all the components so that they are on a LAN within the public cloud region, placement group, and availability zone, you will subject your SAS users to slow WAN speeds that can greatly impact the performance of their SAS applications. A blog, “Does it Matter Where the Various Components of Your SAS Infrastructure Are Installed”, on this subject can be found here: <https://communities.sas.com/t5/Administration-and-Deployment/Does-It-Matter-Where-the-Variou-Components-of-Your-SAS/m-p/483426>.

CLIENTS

The most common placement of SAS clients like SAS® Enterprise Guide®, SAS® Data Integration Studio, and SAS® Studio (even though this is a web-based client, the issue with its placement in comparison to the SAS processing is very important) is within the public cloud infrastructure (either a Windows server or Windows virtual desktops). The question is where to put these Windows systems – in the public cloud in the same availability zone and placement group, on a system within your current data center, or on a desktop/laptop?

The answer depends on the volume of data being transferred back to the SAS client. If there is a lot of data being transferred to populate drop-down windows or to view tables, then having the clients and the backend SAS servers on the same LAN will yield the fastest results. For public clouds, this would mean standing up an instance with a Windows server on it and placing your SAS clients on that instance. Your SAS users would then have to access their SAS clients on the Windows server inside the public cloud infrastructure.

In that case, it is important that the Windows server is located in the same region, placement group, and availability zone as the SAS deployment.

SAS INFRASTRUCTURE

SAS has several tools that allow sharing of SAS data files on-premises with SAS applications that run in a public cloud and vice-versa. While these tools function well, please note that the IO throughput between your data center and the public cloud might be as low as 500 KB/second. Rarely have we seen it more than 20 MB/sec. If the SLAs for your SAS applications in the public cloud can be met with this slow data transfer rate, then it is acceptable to keep your data in a different physical location than where SAS is running.

Our best practice is to have SAS applications and the data associated with them residing in the same data center – whether on-premises or in the same availability zone and placement group in the public cloud.

SOURCE DATA FILES

As we mention above, it is best to have the source data files in the same location as the the systems that will be using these files. If you have to keep your source data files in a different location on the WAN, please note you will have to incorporate the slow WAN performance to the SLA for using SAS in the public cloud. Please note many of the client GUI pull data from the source data files when populating screens, so having the source data files across a slow WAN will greatly impact the performance of the SAS client.

AUTHENTICATION TOOLS

The considerations needed when deciding where to place your authentication tools are similar to those needed when deciding where to place SAS clients or data used by SAS: in the public cloud in the same availability zone and placement group or on a system within your current data center?

This depends on how often your SAS application will interact with your authentication tools for permission to use a file. If this is a high number, it is recommended that you move your authentication tools and associated data into the public cloud. Again, these need to run on a system in the same region, placement group, and availability zone.

AVAILABLE QUICK STARTS FOR SAS

To ease your deployment of SAS, there are several SAS Quick Starts available to you:

SAS 9.4 GRID MANAGER

- Amazon: <https://aws.amazon.com/quickstart/architecture/sas-grid/>
- Azure: not available at the writing of this paper
- Goggle: not available at the writing of this paper

SAS VIYA

- Amazon: <https://aws.amazon.com/quickstart/architecture/sas-viya/>
- Azure: not available at the writing of this paper
- Goggle: not available at the writing of this paper

ARE HIGH AVAILABILITY AND SECURITY REQUIRED?

Depending on your need for high availability (HA), you might require that processes be in place to quickly create a new cloud host instance in the event that one of your existing instances fails. This is more of a failover HA practice. The ability to quickly spin up a new instance is one of the benefits of running SAS in a public cloud.

For SAS Grid Manager customers, please note that a shared file system (for example, Intel Lustre, IBM Spectrum Scale) will remain operational if one of the nodes associated with the shared file system goes down. However, any data that is associated with that node will not be available until the node is restored. This is due to the fact that only one copy of the data is stored by default. It's possible to enable replication services with these shared file systems so that two or three copies of your data are stored on multiple nodes of the shared file system. This does drive up the cost, especially if you have hundreds of terabytes of data.

If your definition of high availability includes disaster recovery, then you will need to look at mirroring your SAS deployment, SAS files, and data store to another region of the public cloud or even to a separate cloud. The SAS application considerations for disaster recovery will apply to the public cloud in the same manner as on-premises infrastructures. Details on what needs to be considered in a Disaster Recovery SAS implementation can be found in [“Do You Have a Disaster Recovery Plan for Your SAS® Infrastructure?”](#) (in *Proceedings of the SAS Global Forum 2016 Conference*).

CONCLUSION

This paper has been written to help SAS customers understand the public cloud instance types that best meet the needs of their SAS workloads. Many of the topics discussed in this paper are based on real-world experiences by SAS customers standing up SAS workloads in public cloud offerings. These customers have learned that there are very few things from a hardware and storage perspective that can be tuned, but the ones that can are listed in this paper. The chief takeaway is that IO throughput is crucial and, unfortunately, is a limiting factor in the success of SAS public cloud deployments.

The intent of this paper is to raise awareness of SAS workload requirements (not only storage but also IO throughput needs) and how best to meet these requirements in a public cloud. The choice of hardware

resources, data stores, and application architecture placement are crucial to achieving the best performance the cloud can offer. The maximum IO throughput to persistent data storage that most public cloud instances can offer is usually much lower than the typical SAS minimum recommended IO throughput of 100 MB/s per core (or more for SAS WORK file systems). You must carefully consider the performance ramifications of moving on-premises SAS workloads to the public cloud, unless you can suffice with a decreased IO bandwidth performance for persistent data activity.

As mentioned in the Introduction, it should be noted that the information discussed in this paper is based on what is available from the public clouds and SAS' experience with the public clouds at the time of the writing of this paper. Public cloud offerings are constantly changing. Therefore, it is in your best interest to understand the rationale used in the selection processes and to consider what was done as a point-in-time design.

REFERENCES

Amazon Cloud. "Amazon EBS-Optimized Instances." Available <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSOptimized.html#ebs-optimization-support>.

Amazon Cloud. "Amazon EC2 Instance Types." Available <https://aws.amazon.com/ec2/instance-types/>.

Google Cloud. "Machine Types." Available <https://cloud.google.com/compute/docs/machine-types>.

Microsoft Azure. "High Performance Compute VM Sizes." Available <https://docs.microsoft.com/en-us/azure/virtual-machines/windows/sizes-hpc>.

Red Hat. "Optimizing SAS on Red Hat Enterprise Linux (RHEL) 6 & 7." Available http://support.sas.com/resources/papers/proceedings11/342794_OptimizingSASonRHEL6and7.pdf.

Crevar, Margaret. April 2016. "Do You Have a Disaster Recovery Plan for Your SAS® Infrastructure." Available <https://support.sas.com/resources/papers/proceedings17/SAS0565-2017.pdf>.

Crevar, Margaret. August 2018. "Does It Matter Where the Various Components of Your SAS Infrastructure are Installed?" BLOG. Available <https://communities.sas.com/t5/Administration-and-Deployment/Does-It-Matter-Where-the-Variou-Components-of-Your-SAS/m-p/483426>

Smith, Ben. April 2018. "Performance and Tuning Considerations on Amazon Web Services with SAS® 9.4 Using IBM Spectrum Scale™". Available <https://support.sas.com/resources/papers/performance-tuning-considerations-amazon-web-services-sas-9-4-ibm-spectrum-scale.pdf>.

Crevar, Margaret. April 2015. "Performance and Tuning Considerations for SAS Grid Manager 9.4 on Amazon Cloud using Intel Cloud Edition for Lustre File System." Available <http://support.sas.com/rnd/scalability/grid/SGMonAWS.pdf>.

ACKNOWLEDGMENTS

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

CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact:

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

Ande Stelk
SAS Institute Inc
Ande.Stelk@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.