

SAS® GLOBAL FORUM 2021

Paper SAS1201-2021

Best Practices When Moving SAS® to Microsoft Azure

Margaret Crevar, SAS Institute Inc.

ABSTRACT

The best practices presented here are lessons learned from SAS customers who have moved to Microsoft Azure. These best practices are being placed in a [single location](#) to share with other SAS customers looking to move their SAS applications to Microsoft Azure.

INTRODUCTION

DISCLAIMER: Public clouds are enhancing their offerings on a regular basis. Information shared in this presentation are the best practices at the time it is written and may not be applicable in 6+ months.

BEFORE YOU START: You need to understand the compute resource (cores, memory, IO throughput and network bandwidth) needs of your SAS applications. The more details you know about your SAS application, the better you can be guided MS Azure infrastructure needed.

After you understand the SAS applications' compute resource needs, you need to define what the success criteria is for moving to the cloud. Many say they want the SAS application to run faster than on-premises. Please note that accomplishing this may be outside the budgetary constraints for our project. So, if there is a limited budget, please understand you may have to adjust user expectations.

Additional information needed before planning your move to the cloud is to understand where all the pieces of the SAS infrastructure will reside – data sources, client interfaces, third party tools. If any of these are located outside the Azure Proximity Placement Group used for the SAS infrastructure, then performance degradations may occur.

WHAT INSTANCE TYPES TO USE

To help you determine the instance type to use for your SAS application, let's discuss the different SAS components.

- For SAS 9.4 these include mid-tier, metadata tier, compute tier, and if SAS Grid – shared file system nodes.
- For SAS Viya 3.5, these include CAS Controller/Workers, Microservices, Postgres/RabbitMQ, SAS Programming Run Time nodes.

This [link](#) brings you to the list of Azure instance types. Read the description carefully to understand what compute resources are available with **each** instance family.

For the Esv3 series, the instance type contains multiple processor models –Broadwell, Skylake or Cascade Lake. Via the Azure portal, you are unable to select the chip set that will be used for the instance. After an instance is instantiated, using the `lscpu` command will list the CPU Model Name for the system.

For SAS Grid compute nodes and CAS Controller/Workers, we recommend these systems all be the same processor type. This will ensure you get **consistent** performance overall. We strongly suggest investing in Azure Unified (a.k.a. Premier) Support.

Please remember that most public cloud instances list CPUs as virtual CPU(s). You need to divide the number of hyperthreaded vCPUs by 2 to get the number of physical cores.

Series	Memory: GiB	Temp storage (SSD) GiB	Max data disks	** Max cached and temp storage throughput: IOPS/MBps (cache size in GiB)	Max uncached disk throughput: IOPS/MBps	Max NICs	Expected Network bandwidth (Mbps)
Edv4 and Edsv4-series	16	75	4	19000/120(50)	3200/48	2	1000
Ev4 and Esv4-series	32	150	8	38500/242(100)	6400/96	2	2000
M-series	64	300	16	77000/485(200)	12800/192	4	4000
Msv2 and Mdsv2 Medium Memory series	128	600	32	154000/968(400)	25600/384	8	8000
Mv2-series	160	750	32	193000/1211(500)	32000/480	8	10000
Constrained vCPUs	256	1200	32	308000/1936(800)	51200/768	8	16000
> Storage optimized	384	1800	32	462000/2904(1200)	76800/1152	8	24000
> GPU - accelerated compute	504	2400	32	615000/3872(1600)	80000/1200	8	30000
> FPGA - accelerated compute							
> High performance compute							
> vCPU quotas							
Azure VMs with no temp disk							
Azure VM sizes naming							

- Review “Temp storage (SSD) GB” and “Max cached and temp storage throughput: IOPS/MBps (cache size in GB)” to see the size and maximum IO throughput of the local, ephemeral disk. For a **Standard_E64ds_v4** instance, the maximum size of the internal SSD that could be used for temporary SAS file systems (SAS WORK/UTILLOC or CAS_DISK_CACHE) is 2400 GB and the maximum IO throughput is 3872 MB/sec (121 MB/sec/physical core). If the size of this ephemeral storage is okay for your SAS WORK/UTILLOC, then you may not need to add Premium Storage to the instance.
- Review “Max uncached disk throughput IOPS/MBps” - **The last line in the above picture is for the Standard_E64ds_v4 instance** (one of the most popular MS Azure instances that is being used for SAS compute systems), the maximum IO throughput (instance total, not per physical core) is 1200 MB per second. For a 32 physical core system, this means 37.5 MB/sec/physical core IO bandwidth for all the data that will be stored on external Premium Storage (used for SAS WORK/UTILLOC). If you need more IO throughput per physical core to the external Premium Storage, you can constrain the number of cores in the instance. There will be more details on “constraining cores” later in this paper.
- Review “Max NICs/Expected network bandwidth (Mbps)” to see what the maximum network egress bandwidth is. For the **Standard_E64ds_v4** instance, the maximum network egress bandwidth is constrained to 30 Gigabit/sec.

WHAT STORAGE TYPE TO USE

There are multiple file systems needed by SAS. Gather a list of the ones currently used by your SAS application before reviewing what is needed in Azure.

Permanent SAS Data Storage: This data storage persists between SAS sessions and after reboots of the Azure instances. It can be local file systems or shared file systems (more details in the next section). The [link](#) in this section will point you to the different Premium storage drives to use for locally attached file systems. Like the instance types, there is a maximum IO throughput per Premium Disk.

Temporary SAS Data Storage: This data storage does not persist between SAS sessions or after reboots of the Azure instance. Ephemeral storage can be used for this storage. This storage tends to need the fastest IO throughput.

What IO throughput is needed for each: You need to gather this from your existing SAS infrastructure and determine if storage in Azure can match what you have on-premises or if you need faster IO throughput.

To get consistent instance-to-instance IO and throughput, ensure all your instances are in the same [Azure Proximity Placement Group](#).

As discussed on the Azure Instance section, there is a maximum IO throughput to attached storage. If you need more IO throughput, then you can *Utilize Constrained Cores* with Azure instances to reduce the number of physical cores presented to the instance's operating system, thereby increasing the IO bandwidth per physical core.

As a reminder, SAS temporary files and directories such as SAS WORK, SAS UTILLOC and CAS_DISK_CACHE should be placed on storage with the highest proven throughput possible. Today that usually means Premium Storage or the instance's local ephemeral SSD.

WHICH SHARED FILE SYSTEM TO USE

A shared file system is required if you are using SAS Grid. As with everything, there are pros and cons for which shared file system to use.

- **EXASCALER CLOUD SHARED FILE SYSTEM:** Please review this blog on how to use EXAScaler Cloud (Lustre) with SAS. [EXAScaler Cloud by DDN: A shared file system to us... - SAS Support Communities](#). EXAScaler Cloud has replaced the community version of Lustre.
- **IBM SPECTRUM SCALE:** Please review this blog on how to use IBM Spectrum Scale with SAS. [Sycomp Storage Fueled by IBM Spectrum Scale: A new... - SAS Support Communities](#)
- **AZURE NETAPP FILES (ANF):** Please review this blog on how to use ANF with SAS. [Azure NetApp Files: A shared file system to use wi... - SAS Support Communities](#) Please note that there are maximum number of physical cores that can be used with ANF storage. These are detailed in the blog. And ANF should not be used for any file system that will support heavy write workloads (including SAS WORK/UTILLOC and CAS_DISK_CACHE)

TUNING GUIDELINES

These guidelines are from "lessons learned" by the initial SAS in the Azure cloud.

1. To get consistent instance-to-instance IO and network throughput, ensure all your instances are in the same Azure Proximity Placement Group.

2. Avoid sporadic NMI lockups that might hold processing while a thread waits for an available vCPU when using RHEL 7.x (3.10 kernel) with SAS compute nodes. There is a known issue in the iSCSI and SCSI drivers in this kernel which can cause CPU lock ups when under heavy IO load.
3. To achieve optimal network bandwidth, Azure Accelerated Networking must be enabled. Accelerated Networking is available on any Azure VM with 2 or more physical cores.
4. In addition to Accelerated Networking, SAS needs to be on an isolated cloud VNET, Resource Group, etc. This VNET should “share nothing” with other customer infrastructures. The exception is placing the instances for your shared file system and RDBMSs dedicated to SAS on this VNET as well.
5. Azure Default VM Network MTU Size - Azure strongly recommends the default network MTU size of 1500 not be adjusted because Azure’s Virtual Network Stack will attempt to fragment a packet at 1400 bytes. To learn more, please review this [“Azure and MTU”](#) article.

More details on the above can be found in this blog post: [Best Practices for Using Microsoft Azure with SAS® - SAS Support Communities](#)

SAMPLE AZURE INSTANCE

To summarize the above, the following is a good example configuration for SAS 9.4 or SAS Viya 3.5 compute nodes.

STANDARD_E64-32DS_V4: specs for this system:

- Cascade Lake processor.
- 16 physical cores (32 vCPUs)
- 504 GB RAM
- The internal 2,400 GB SSD drive can be used for SAS temporary file systems, but it cannot be increased in size. The throughput for this storage is 3,872 MB/sec which equates to 121 MB/sec/physical core. SAS recommends at least 150 MB/sec/physical core.
- 30 Gigabit egress network connectivity
- For persistent storage, use six P30 Premium Disks striped together for a total of 6 TBs. If more disk space is needed, then add more P30 disks or larger Premium Disks. Remember the maximum IO bandwidth to the E64 instance is 1,200 MB /sec. With the constrained cores, this equates to 75 MB/sec/physical core for the Standard_E64-32ds_v4. SAS recommends at least 100 MB/sec/physical core.

CONCLUSION

There are many resources, configuration settings and constraints to check within Azure to configure an instance to meet the needs of your SAS application.

It is possible you may have to provision an instance with more physical cores (with or without a constrained core count) in order to get the IO throughput required by your SAS application. Likewise, you may also have to over-provision storage capacity to achieve the IO throughput needs for your SAS application.

As always, there are cost versus performance choices. These selections need to be based on your SLAs and business needs for SAS applications running in Azure. Please remember things are always changing in the public cloud.

Lastly, please join the SAS Communities for “Administration” so you can be kept apprised of new information related to running SAS in the public cloud.

REFERENCES

Crevar, Margaret. 2021. [Best Practices for Using Microsoft Azure with SAS® - SAS Support Communities](#)

Microsoft Docs. 2021 [VM sizes - Azure Virtual Machines](#)

Microsoft Docs. 2021 [Azure Premium Storage: Design for high performance - Azure Virtual Machines](#)

Crevar, Margaret. 2021. [EXAScaler Cloud by DDN: A shared file system to us... - SAS Support Communities](#)

Crevar, Margaret. 2021. [Sycomp Storage Fueled by IBM Spectrum Scale: A new... - SAS Support Communities](#)

Crevar, Margaret. 2021. [Azure NetApp Files: A shared file system to use wi... - SAS Support Communities](#)

ACKNOWLEDGEMENTS

Many thanks to SAS R&D, SAS Technical Support, Microsoft Azure, Azure NetApp Files, Sycomp and DDN experts for reviewing this post.

RECOMMENDED READING

Marson, Barry. 2019. [Optimizing SAS on RHEL 6 and 7 \(redhat.com\)](#)

Crevar, Margaret. 2020. [Important Performance Considerations when Moving SAS to the Public Cloud](#)

CONTACT INFORMATION

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

Margaret Crevar
SAS
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.