



How Do I Move SAS[®] Applications to a Public Cloud?

Margaret Crevar





Margaret Crevar **Senior Manager, Software Development, SAS**

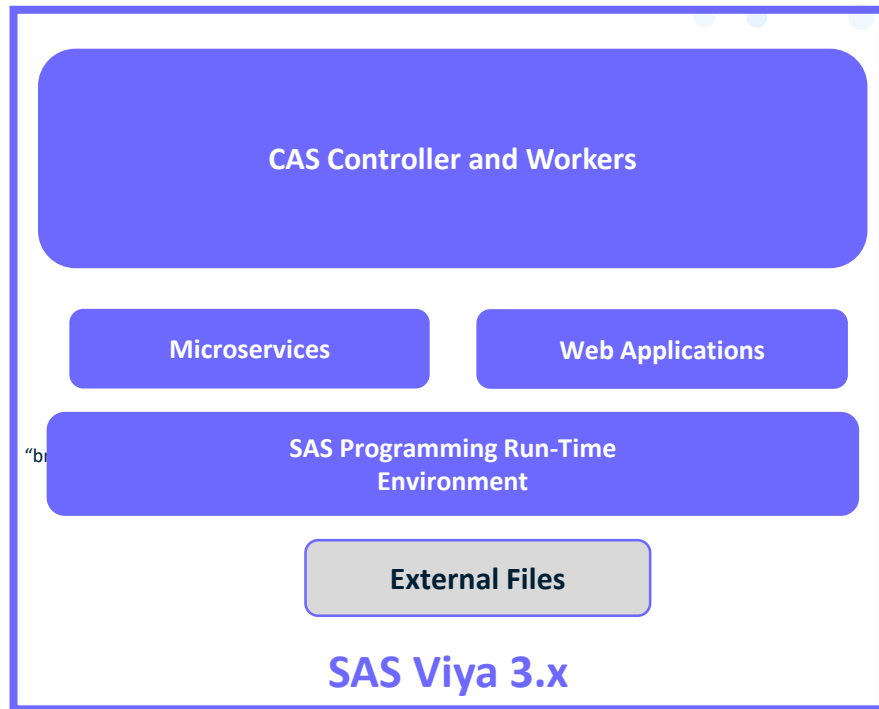
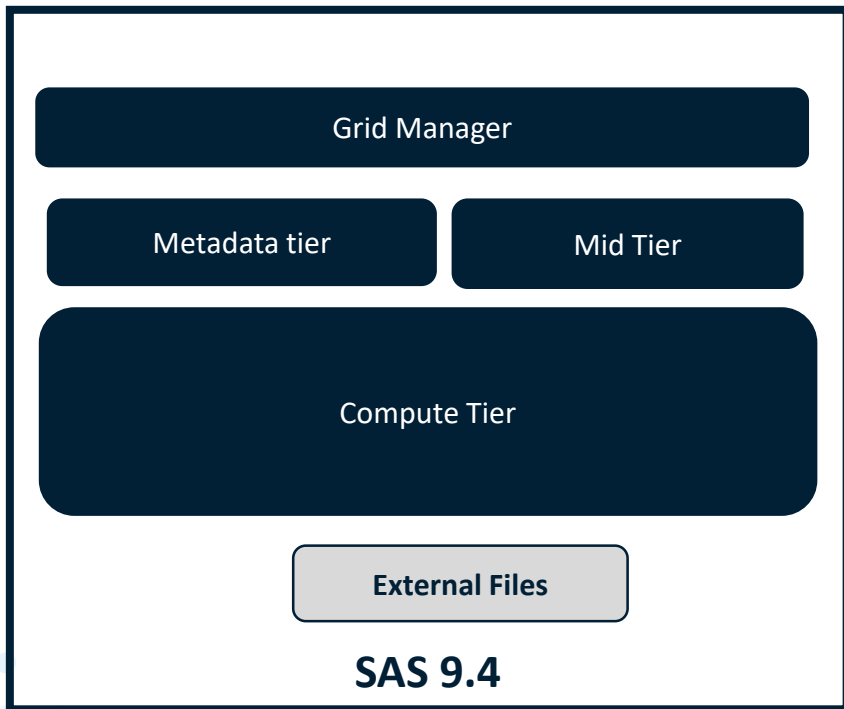
Crevar's duties include helping SAS customers with performance issues. She uses experience to teach new SAS customers to correctly configure their hardware and storage. Her goal is to ensure optimal performance for SAS at customer sites. Crevar works closely with SAS partners to produce tuning guidelines and best practices.

Disclaimer

- 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 clouds are always enhancing their offerings
- SAS experience in the public clouds is growing on a daily basis.



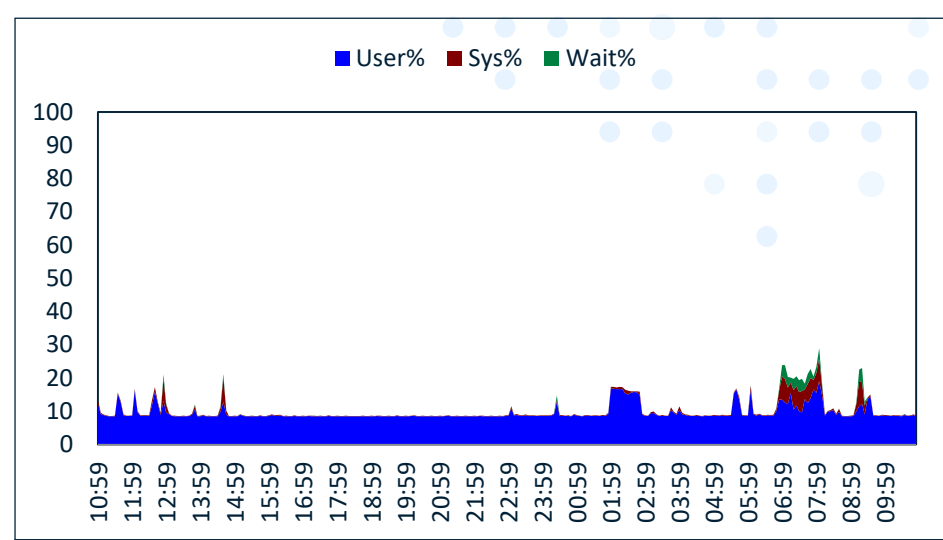
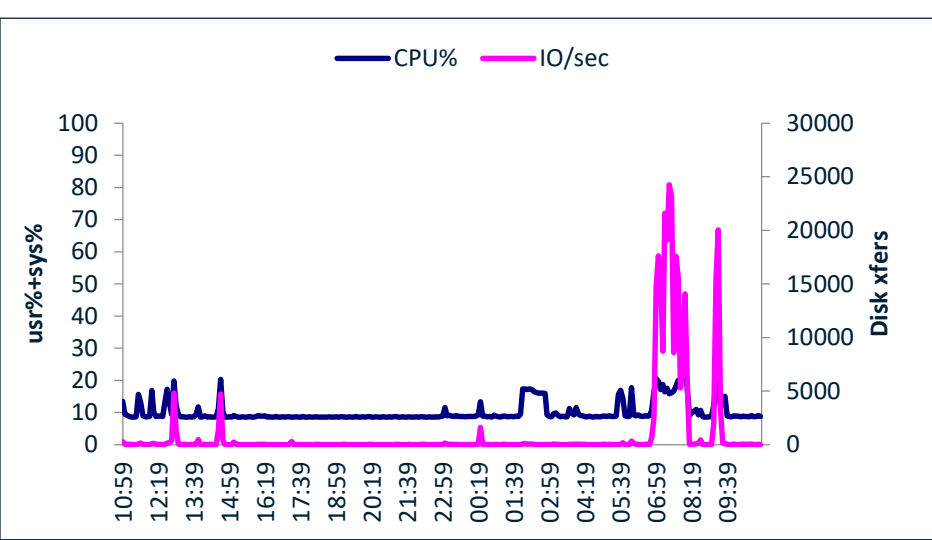
Before you Start

- A good understanding of
 - the SAS application workload and their SLAs
 - all layers and components of the SAS infrastructure
- Definition of success criteria
- Compute resources being used by existing SAS application



Compute Resources Being Used

- CPU cycles – spikes and average
- Memory – spikes and average
- IO throughput
 - Persistent files
 - Temporary files
- Network bandwidth



Agenda

- What instance type to use
- How to use the correct storage type
- Which shared file system to use when deploying SAS Grid Manger
- Where to place SAS components

What Instance Type to Use

- SAS 9 Servers
 - SAS Compute Tier
 - Shared File System Storage Required for SAS Grid
 - SAS Mid-tier and Metadata Servers
- SAS Viya Servers
 - CAS Controller and Worker Nodes
 - MicroServices and Postgress/RabittMQ Node(s)
 - SAS Programming Run Time Node(s)

How to Use the Correct Storage Type

- Permanent SAS Data Storage
- Temporary SAS Data Storage
- What IO throughput is needed for each area

AWS Instance for SAS Compute Tier

- i3en-6xlarge
 - 12 physical cores – 24 vCPUs
 - 192GB of RAM
 - Two 7,500 NVMe drives. Striped together, we get over 200 MB/sec/physical core. This is ephemeral space which is good for temporary SAS files.
 - Persistent storage:
 - The maximum IO throughput is based on available network.
 - EBS ST1 storage
 - 25 Gigabit network

Azure Instance for SAS Compute Tier

- Standard_E32ds_v4
 - 16 physical cores – 32 vCPUs
 - 256GB of RAM
 - 1200GB internal SSD drive. Maximum IO throughput is 1936 MB/sec with equals 121 MB/sec/physical core. This is ephemeral space which is good for temporary SAS files.
 - Persistent storage:
 - The maximum IO throughput is 768 MB/second which equals 48 MB/sec/physical core.
 - Premium storage will need to be multiple Premium Disks stripe together. Which one to use will depend on the total capacity needed persistent storage. Four P30s = 4GB of space or Three P40s = 6GBs of space.
 - 16 Gigabit network
 - May have to use constrained cores to get required IO throughput.

Constrained Azure Instance for SAS Compute Tier

- Standard_E32-16ds_v4
 - 8 physical cores – 16 vCPUs
 - 1200GB internal SSD drive. Maximum IO throughput is 1936 MB/sec with equals 242 MB/sec/physical core. This is ephemeral space which is good for temporary SAS files.
 - Persistent storage:
 - The maximum IO throughput is 768 MB/second which equals 96 MB/sec/physical core.
 - Premium storage will need to be multiple Premium Disks stripe together. Which one to use will depend on the total capacity needed persistent storage. Four P30s = 4GB of space or Three P40s = 6GBs of space.
 - 16 Gigabit network

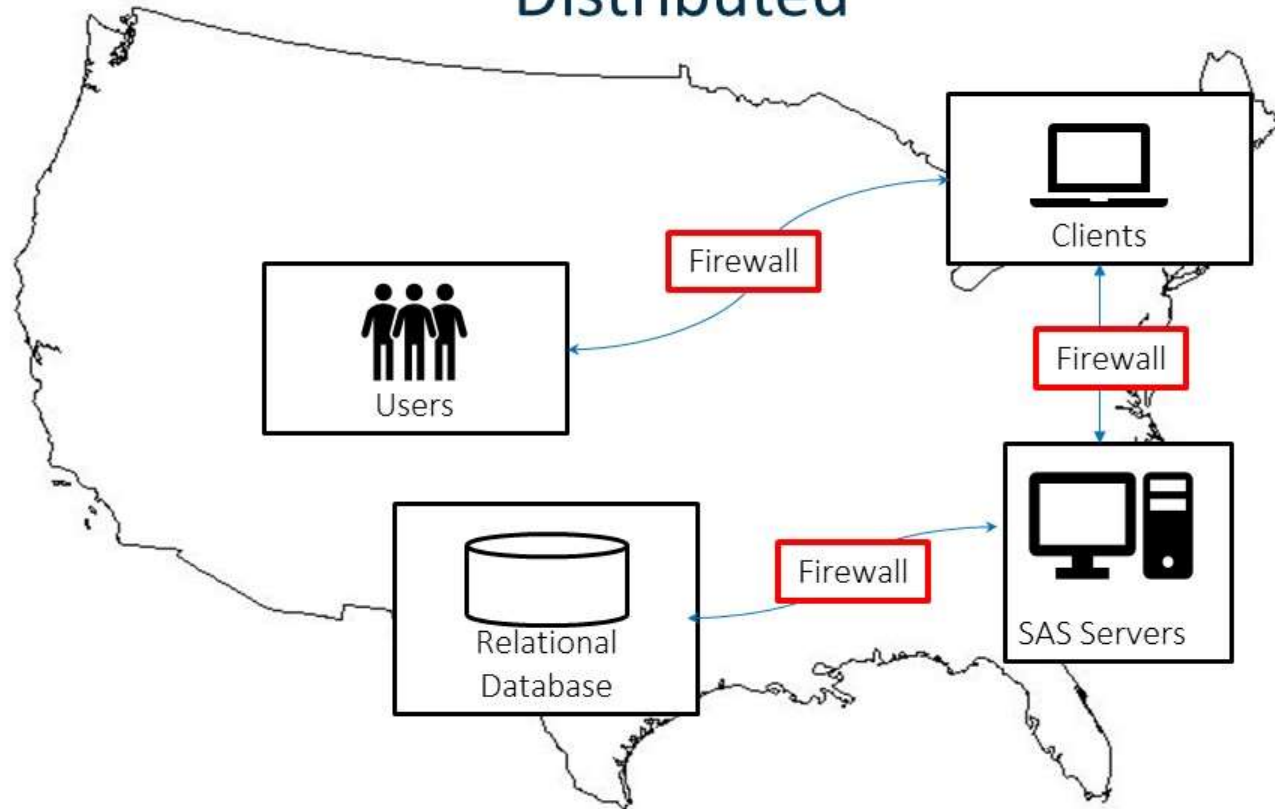
Which Shared File Systems to Use

- DDN's Lustre
- IBM Spectrum Scale (AWS only)
- Amazon FSx for Lustre (AWS only)
- Amazon EFS Storage (AWS only)
- Azure NetApp File (Azure only)
- NFS

Where to Place SAS Components

- SAS clients
- Data to be used by SAS
- Third Party tools
 - Security
 - Authentication
- High Availability

Architecture Distributed



Conclusion

- Need to understand the complete workload and expected SLAs of the workload
- Need to know that IO throughput is limited to storage external to the instance
- Need to limit access to data outside of public cloud data center
- Instances, storage types, ... are always changing in the public clouds
- Information learned will be posted to
 - support.sas.com/kb/62/239.html
 - communities.sas.com/t5/Administration-and-Deployment/bd-p/sas_admin
 - communities.sas.com/t5/Architecture/bd-p/architecture

Recommended Reading

- All public clouds
 - www.sas.com/content/dam/SAS/support/en/sas-global-forum-proceedings/2020/4312-2020.pdf
- AWS
 - aws.amazon.com/ec2/instance-types/
 - docs.aws.amazon.com/whitepapers/latest/aws-overview/storage-services.html
 - d1.awsstatic.com/whitepapers/best-practices-for-deploying-sas-server-on-aws.pdf?did=wp_card&trk=wp_card
- Azure
 - docs.microsoft.com/en-us/azure/virtual-machines/sizes-storage
 - docs.microsoft.com/en-us/azure/virtual-machines/windows/premium-storage-performance
 - communities.sas.com/t5/Administration-and-Deployment/Best-Practices-for-Using-Microsoft-Azure-with-SAS/m-p/676833#M19680
- Google
 - cloud.google.com/compute/docs/machine-types
 - cloud.google.com/compute/docs/disks/performance#c2-disk-limits

Q & A

Explore Helpful User Resources

[Ask the Expert](#)

View other user webinars that provide insights into using SAS products to make your job easier.

[FREE Training](#)

Learn from home – free for 30 days. Get software labs to practice and online support if needed.

[SAS Support Communities](#)

Ask questions, get answers and share insights with SAS users.

[SAS Analytics Explorers](#)

An exclusive platform to collaborate, learn and share your expertise. Gain access to a diverse network to advance your career. Special rewards and recognition exclusively for SAS users.

[SAS Users YouTube Channel](#)

A plethora of videos on hundreds of topics, just for SAS users.

[Newsletters](#)

Get the latest SAS news plus tips, tricks and more.

[Users Groups](#)

Meet local SAS users, network and exchange ideas – virtually.

[SAS Profile](#)

If you haven't already done so, create your SAS Profile to access free training, SAS Support Communities, technical support, software downloads, newsletters and more.