



SAS on Azure

Experiences from the field

Roeland Nieuwenhuis

Principal Cloud Solution Architect
Customer Success CAT

Choices ...

1. Pick the right instance type
2. Pick the right storage
3. Avoid common problems

Success!

Instance types: things to be aware of

- SAS uses MKL for some of its analytics, be cautious using AMD
- SAS refers to physical cores, Azure to vCPUs. Ratio is 1:2.
- Go with Linux. Avoid golden images and run latest of RHEL
- Always use Accelerated Networking

Instance types: what can be considered

- Esv3
 - Not recommended, older generation, you can get different CPU generations too
- Esv4 and Esv5
 - Good fit, but no local storage and slower remote disk
- Edsv4
 - Recommended VM for SAS Viya workloads right now
- Ebsv4
 - Recommended VM for SAS 9 Grid, developed with SAS workloads in mind

Instance types: not a great fit

- D and F series
 - Insufficient memory
- H series, N series
 - No accelerated networking, only fits when using RDMA/IB
 - SAS doesn't seem to use GPUs (yet?)
- L series
 - Can work for specific calculations (e.g., Monte Carlo)
 - Lsv2 ran into trouble with network throughput due to NUMA
 - Intel MKL based analytics are heavily penalized
 - Next version of L series is coming soon and fixes these issues

Storage choice

- Need to cater for two storage systems
 - Ephemeral (SASWORK, CAS_DISK_CACHE, SAS UTILLOC) as fast as possible – 75MB/s per vCPU
 - Shared file system (SASDATA, CASDATA) – next slides
- Use local storage whenever possible, premium SSDs as a last resort
 - Use lvm striped, count = 2, 64KB block size on ext4 or xfs
 - Ultra SSDs are not useful, the workload is not IOPS heavy and it'll hit VM cap first
- You can use network offloads for ephemeral, but not recommended, especially not for Viya as it is write latency sensitive

Storage choices: Azure NetApp Files

- Good fit for smaller workloads (48 cores max) or workloads that have limited data (i.e., a few TB)
- Be aware of throughput max. of 1.8GB/s write and 4.5GB/s read
 - Will be addressed in the near future
- Currently caps out at 100TB of data per volume
 - Can stack volumes, but gets complicated

Storage choices: Lustre

- Linearly scaling throughput system, capable of getting to 100+GB/s of throughput for very large file systems (PB+)
- Doesn't offer much redundancy, data loss protection or encryption (no "enterprise features") out of the box
- Common choices are:
 - DDN EXAScaler (former WhamCloud) or CycleCloud
 - Use Terraform templates for Lustre + Lemur on GitHub

Storage choices: Sycomp GPFS

- IBM Spectrum Scale managed through marketplace offer
- Like Lustre, provides linearly scaling filesystem capable of 100+GB/s throughput at low latency
- Has encryption (SKLM), encryption end-to-end, backup strategies and more already built in

Understand your authentication requirements

- Make sure you know where you want to use Azure AD and where Kerberos/AD
- Avoid mixing B2B and B2E + Kerberos
- If you need to use NFSv4.1 + ACLs, stay with Kerberos
- Do not use AADDS

Must do's

- Deploy a supported OS, RHEL 7.9 is the default and preferred. Tune kernel if needed to avoid dumps.
- Deploy all resources forced to a zone, single VNet, multiple subnets and use a PPG
 - Avoid VNet peering to data sources, unless you want to spend lots of \$\$ on peering traffic
 - Use Accelerated networking
- Monitor and evaluate shared storage utilization and throughput and right size it, run `rhel_io` to test storage sanity
- Check reports for machine utilization, especially disk and network to see if you are being constrained, resize if needed
- Snooze clusters when using SAS Viya 4 on AKS

Avoid (issues we've seen)

- Azure Disk Encryption
 - Penalizes your performance
 - Use BYOK with SSE instead
- Using ILBs in front of mid tiers
 - Do not use ILBs in front of SAS Mid tier, as it does a self reference: not supported. Use App Gateway instead.
- `MKL_DEBUG_CPU_TYPE=5` – no longer works
- Do not touch MTUs or you'll get fragmentation

“This is great documentation; I wish I saw it before!”

Customers after seeing our aka.ms/sasdocs



Thank you

