Best Practices When Moving SAS to the MS Azure Cloud

Margaret Crevar

SAS Institute

Copyright © SAS Institute Inc. All rights reserved.

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 daily.

Copyright © SAS Institute Inc. All rights reserved.

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

Location of

- source data
- third-party tools
- SAS clients

Copyright © SAS Institute Inc. All rights reserved.

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 Postgres/RabbitMQ Node(s)
- SAS Programming Run Time Node(s)

Edv4 and Edsv4-series - Azure Virtual Machines | Microsoft Docs

- Maximum uncached disk throughput IOPS/MBps
- Maximum NICs/Expected Network bandwidth (Mbps)
- Temp storage (SSD) GB

Copyright © SAS Institute Inc. All rights reserved.

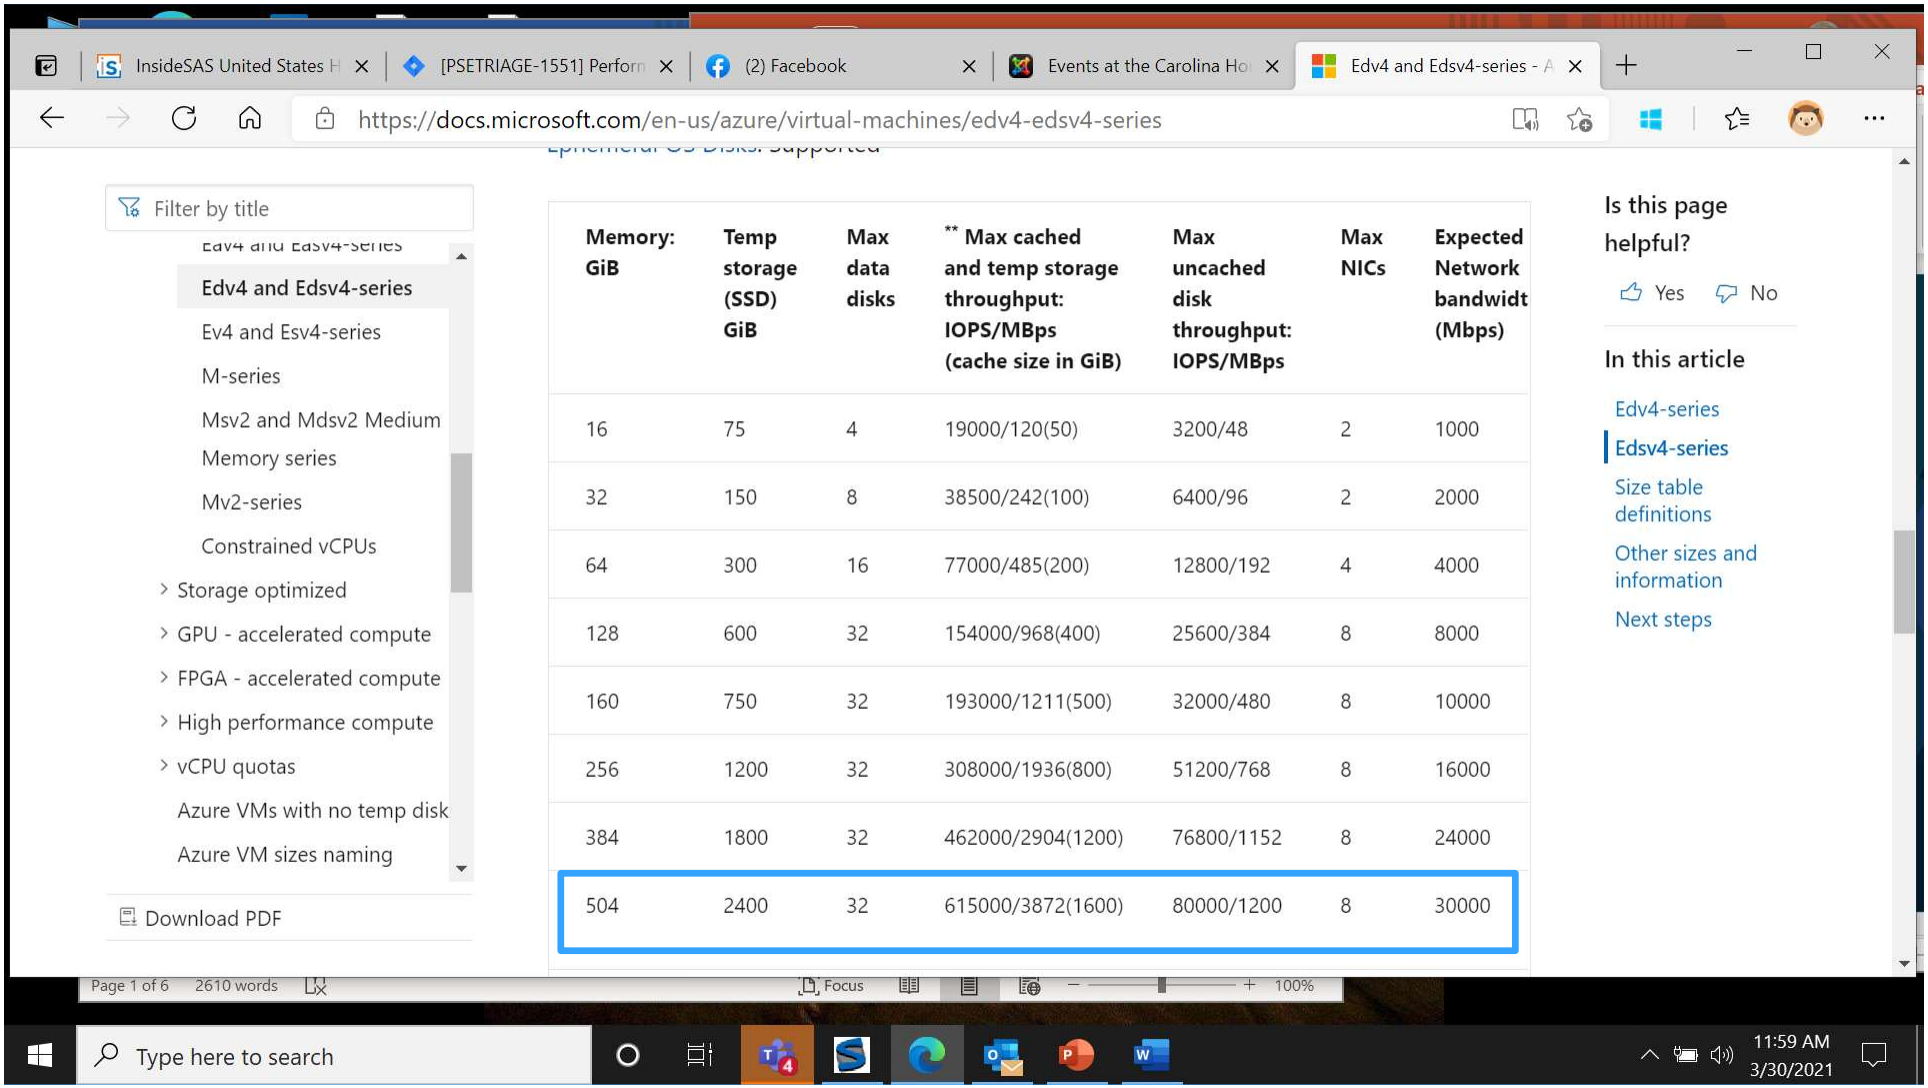
InsideSAS United States H
[PSETRIAGE-1551] Perform
(2) Facebook
Events at the Carolina Ho
Edv4 and Edsv4-series - A
https://docs.microsoft.com/en-us/azure/virtual-machines/edv4-edsv4-series
Filter by title
Edv4 and Edsv4-series
Ev4 and Esv4-series
M-series
Msv2 and Mdsv2 Medium Memory series
Mv2-series
Constrained vCPUs
Storage optimized
GPU - accelerated compute
FPGA - accelerated compute
High performance compute
vCPU quotas
Azure VMs with no temp disk
Azure VM sizes naming
Download PDF
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 bandwidt (Mbps)
16 75 4 19000/120(50) 3200/48 2 1000
32 150 8 38500/242(100) 6400/96 2 2000
64 300 16 77000/485(200) 12800/192 4 4000
128 600 32 154000/968(400) 25600/384 8 8000
160 750 32 193000/1211(500) 32000/480 8 10000
256 1200 32 308000/1936(800) 51200/768 8 16000
384 1800 32 462000/2904(1200) 76800/1152 8 24000
504 2400 32 615000/3872(1600) 80000/1200 8 30000
Is this page helpful?
Yes
No
In this article
Edv4-series
Edsv4-series
Size table definitions
Other sizes and information
Next steps
Page 1 of 6 2610 words
Focus
100%
Type here to search
11:59 AM
3/30/2021

What Storage Type to Use

Permanent SAS Data Storage

Temporary SAS Data Storage

What IO throughput is needed for each area

https://docs.microsoft.com/en-us/azure/virtual-machines/windows/premium-storage-performance

- Premium Storage
- Temp storage (SSD) GB
- Azure Proximity Placement Group
- Utilize Constrained Cores

Copyright © SAS Institute Inc. All rights reserved.

Which Shared File Systems to Use

- DDN’s EXAScaler (Lustre)
- IBM Spectrum Scale (GPFS)
- Azure NetApp File (NFS)

Copyright © SAS Institute Inc. All rights reserved.

Tuning Guidelines

- Azure Proximity Placement Group
- Avoid sporadic NMI lockups
- Azure Accelerated Networking
- Use isolated VNET, using private IP addresses and private DNS.
- Azure Default VM Network MTU size
- Best Practices for Using Microsoft Azure with SAS® - SAS Support Communities

Copyright © SAS Institute Inc. All rights reserved.

Sample Azure Instance

Standard_E64-32ds_v4 instances (Cascade Lake)

- 16 physical cores
- 504 GB RAM
- 2,400 GB internal SSD drive (ephemeral)
- Maximum IO throughput to Premium Storage 1,200 MB/sec
- Expected Network bandwidth 30,000 Mb/sec

Locally attached Premium Storage (non-SFS)

- Use six P30 drives
- Each drive is 1 TB in size
- 200 MB per second per drive

Copyright © SAS Institute Inc. All rights reserved.

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

Copyright © SAS Institute Inc. All rights reserved.

Questions?

sas.com

Copyright © SAS Institute Inc. All rights reserved.