SOLUTION BRIEF
OCTOBER 2020

HIGHLIGHTS

- TAP INTO THE POWER OF AMD EPYC PROCESSORS TO ACCELERATE DATABASES AND APPLICATIONS
- CHOOSE FROM PRECONFIGURED IMAGES WITH WINDOWS SERVER, RED HAT ENTERPRISE LINUX, SUSE ENTERPRISE LINUX SERVER, OR UBUNTU LINUX
- MIGRATE APPLICATIONS CURRENTLY RUNNING ON OTHER CLOUD INSTANCES TO AMD EPYC PROCESSOR-BASED INSTANCES WITH LITTLE TO NO MODIFICATION
- OPTIMIZE YOUR DEPLOYMENTS USING BEST PRACTICES FROM AMD AND MICROSOFT

AMD EPYC™ Processors Power Preconfigured Microsoft SQL Server® Virtual Machines in the Cloud

You get your choice of infrastructure when deploying in the cloud. Now Microsoft Azure gives you access to new levels of horsepower to accelerate your workloads: preconfigured Microsoft SQL Server on Azure Virtual Machines (Azure SQL VMs) instances powered by AMD EPYC processors.

INTRODUCTION

Companies rely on many types of applications to manage the business and make better decisions—data warehousing, online transaction processing (OLTP). online analytic processing (OLAP), reporting, and more. As the number of applications and data volumes grow, database sprawl and the complexity it adds are a common IT concern. That's why many IT managers use public clouds for their database and analytic application deployments. Microsoft Azure SQL VM instances, powered by AMD EPYC™ processors, give you the flexibility to deploy in the cloud without sacrificing performance. Cloud deployments provide the opportunity to:

- Tap into cloud resources
- Accelerate workloads and insight
- Simplify your data center infrastructure and operations
- Help reduce both hardware and license costs

MICROSOFT AZURE SQL VM INSTANCES POWERED BY AMD EPYC PROCESSORS

Using SQL Server on Azure Virtual Machines allows you to use full versions of Microsoft SQL Server without having to manage any server-based, on-premises hardware. During deployment, you choose the underlying compute and storage infrastructure you prefer based on the performance and cost structure that fits your needs and budget.

You can configure your Azure SQL VM instances to accelerate real-world database and application performance. These Da and Ea series instances seamlessly support a preconfigured software stack so that you can migrate your applications with little or no modification. Available for Microsoft Windows® and Linux®, these instances are preconfigured for general-purpose (Da) and memory-intensive (Ea) workloads. Both types of instances, powered by AMD EPYC™ 7452 processors, deliver high throughput, low latency, and directly mapped local temporary storage to Azure SQL VM deployments.

AMD EPYC Processors Power Microsoft SQL Virtual Machines in the Cloud

- **Da SERIES:** provides a balanced CPU-to-memory ratio to support general-purpose workloads such as testing and development, small to medium databases, applications, and application servers. This series allows Azure SQL VM to use AMD EPYC 7452 processors and support up to 96 virtual CPUS (vCPUs), 384 GB of memory, and 2,400 GB of SSD-based temporary storage. By default, 8 MB of the L3 cache is dedicated to every 8 cores.
- **Ea SERIES:** provides a high memory-to-core ratio, making it ideal for relational database servers, workloads that benefit from medium to large caches, and in-memory analytics. This series allows Azure SQL VM to use AMD EPYC 7452 processors and support up to 96 vCPUs, 672 GB of memory, and 2,400 GB of SSD-based temporary storage.

PERFORMANCE YOU CAN COUNT ON

Across virtually all workloads, AMD EPYC processors set the new standard for cloud deployments and deliver breakthrough workload performance. AMD engineers conducted tests derived from industry standard benchmarks on Azure SQL VMs with a variety of Ea instance types to demonstrate the performance characteristics you can expect.[1] Three instance types were tested with different vCPU and memory configurations using the same OLTP benchmark workload derived from the TPC-C™ benchmark.[2] Five test runs were performed and the median result selected.

The workload models an order fulfillment system. The database receives requests for data, adds new data, and makes multiple changes to the data from many users. The results show sustained transaction throughput and predictable scaling between E8as, E16as, and E32as instances on the Windows Server 2016 operating system. With this level of optimized I/O throughput for database applications, you can size your Microsoft SQL VM for current needs and scale the system to support increasing demand while maintaining consistent, predictable performance.

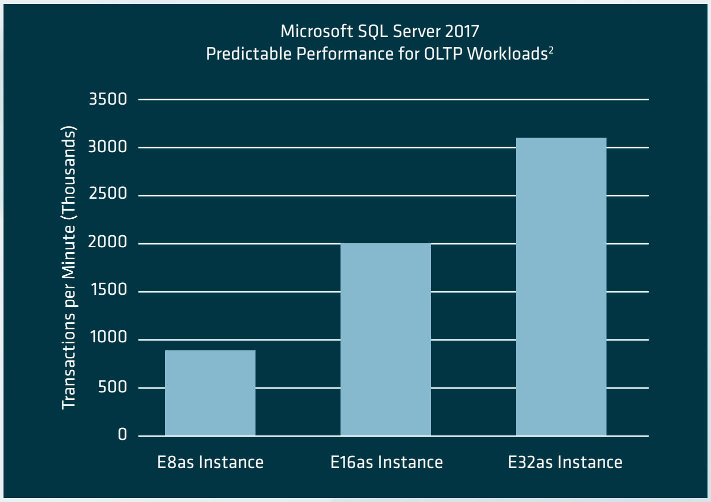

BEST PRACTICES FOR DEPLOYMENT

When deploying Azure SQL VMs, start by using the database performance-tuning options that are used in your on-premises deployment. These options may need modification as the performance of a relational database in a public cloud depends on many factors, such as the size of a virtual machine and data disk configuration. Azure SQL VMs provisioned in the Azure portal follow general storage configuration best practices. After provisioning, consider applying other optimizations described below. Base your choices on your workload and verify performance characteristics through testing.

VIRTUAL MACHINE PERFORMANCE

Getting the most out of your cloud deployment requires tuning the Azure SQL VM configuration. The following suggestions can help you optimize your deployment for a given workload.

- Start by collecting the CPU, memory, and storage throughput requirements of the workload at peak times. Use built-in performance counters to collect read and write I/O operations per second (IOPS) requirements and storage throughput needs for data, log, and temp database files.
- Understand both the throughput and IOPS requirements of the workload as different Azure SQL VMs configurations have different scale limits for the factors.

AMD EPYC Processors Power Microsoft SQL Virtual Machines in the Cloud

- Add 20% more throughput (IOPS) capacity than the workload requires to provide headroom.
- Evaluate infrastructure instance sizes that support the requirements identified and select the best fit.

DATABASE CONSIDERATIONS

The performance of a relational database in a public cloud depends on many factors, such as the size of the virtual machine and memory optimizations. The best practices outlined below focus on achieving the best performance for Azure SQL VMs using AMD EPYC processors. Less demanding workloads may require fewer optimizations. Consider your performance needs, costs, and workload patterns as you evaluate these recommendations.

DATABASE CHECKPOINTS

Databases write data asynchronously to speed performance. Modifications to data pages are done in memory (in the buffer cache), with checkpoints issued periodically to flush modified pages and transaction log data to persistent storage. After an unexpected shutdown or crash, the database starts from the last checkpoint and applies changes contained in the log during the recovery process. The time interval between checkpoints depends on whether the workload is write- or read-intensive.

The recovery interval setting largely dictates the recovery time. The higher the recovery interval setting, the longer the time between checkpoints and the longer the time to recover the database after an event. Although the default settings typically provide optimal recovery performance, modifying the recovery interval might improve performance if the recovery process takes significantly longer than 1 minute when long-running transactions are not being rolled back, or when frequent checkpoints negatively affect database performance.

If you decide to increase the recovery interval setting, increase it by small increments and evaluate the effect of each change on recovery performance. Keep in mind that as the recovery interval setting increases, database recovery takes that many times longer to complete. For example, if you change the recovery interval to 10 minutes, the recovery process takes approximately 10 times longer to complete than when the recovery interval is set to 1 minute. For more information on checkpoint options and settings, read Database Checkpoints (SQL Server).*

TRACE FLAGS

Trace flags provide a way tune the database to support workloads that demand the utmost performance. Trace flags tend to be helpful for Microsoft SQL VMs with at least 8 vCPUs, at least 8 GB of memory, and those experiencing high I/O rates (more than 10,000 physical IOPS or more than 500 MB/second). The following table highlights the key trace flags to consider. See Tuning options for SQL Server when running in high performance workloads* for more information.

LOCK PAGES IN MEMORY

Sometimes Microsoft SQL Server pages memory to virtual disks to free memory for operation. If database performance decreases suddenly, the server stops responding for a short time, or applications experience a timeout trying to connect to the database, it may be due to paging delays. If you suspect that memory is running low and database pages are being paged to disk, you can set a policy to force pages to be kept in physical memory and help improve performance. See Enable the Lock Pages in Memory Option (Windows)* for the steps to take to lock pages in memory.

DATABASE I/O

When running Azure SQL VMs, we recommend that you:

- Parallelize application requests to premium SSDs to achieve best results.
- Use database compression with I/O-intensive workloads.
- Use instant file initialization to reduce the time needed for initial file allocation.
- Keep (or move) all databases, error logs, and trace files to data disks.

AMD EPYC Processors Power Microsoft SQL Virtual Machines in the Cloud

STORAGE CONSIDERATIONS

Azure SQL VMs define and use several types of disk structures to operate. Each disk type is designed for specific constructs and has its own set of best practices to help ensure optimal database and workload performance.

- **OS DISK:** During provisioning, Microsoft Azure attaches an OS disk to the Azure SQL VM to be used as the operating system disk. This virtual drive can be booted and mounted as a running operating system. Avoid using the OS disk for database storage or logging. For performance-sensitive workloads, use data disks instead of the OS disk.
- **TEMPORARY DISK:** Each Microsoft SQL VM contains a temporary disk that is used for scratch space. Because this disk is not persistent, do not store user database files or user transaction logs on the temporary disk. For mission-critical workloads, place TempDB on the local SSD. If you use premium SSDs with your virtual machine, you can also store TempDB on a premium SSDs with read caching enabled.
- **DATA DISKS:** Additional disks can be added to your configuration as data disks. These data disks (Premium SSDs and Ultra Disks) are stored in storage as page blobs.

  Premium SSDs offer the best price/performance and are ideal for data and database log files.[3] Configure a ReadOnly cache for data files and no cache for the log file. In highly available configurations, these premium file shares can also be used for the database failover cluster instance.

  Ultra Disks are ideal for workloads that require less than 1 ms storage latencies. If low latencies are only required for the log file and not for data files, provision the Ultra Disk at required IOPS and throughput levels only for the log file.

Other storage best practices include:

- **USE DISK STRIPING TO INCREASE THROUGHPUT.** To determine the number of data disks, analyze the number of IOPS and bandwidth required for your log file(s), and for your data and TempDB file(s). Note that different Da and Ea instances have different limits on the number of IOPS and bandwidth supported.
- **USE THE RIGHT CACHING POLICY FOR YOUR STORAGE CONFIGURATION.** Caching should only be enabled on data disks that use premium SSDs. If separate disks are used for data and log files, enable read caching on the data disks hosting your data files and TempDB data files to accelerate read performance. Do not enable caching on the disk holding the log file to avoid degrading performance.

  When striping disks in a single storage pool, read caching can increase workload performance. If separate storage pools are used for log and data files, enable read caching only on the storage pool for the data files. Some write-intensive workloads may experience better performance with no caching. Test your workloads to determine how they respond to different caching settings.
- **SET THE NTFS ALLOCATION UNIT SIZE APPROPRIATELY.** Use a 64-KB allocation unit size for data files, log files, and TempDB. If TempDB is placed on the temporary disk, the performance gained by using a fast, dedicated disk minimizes the need for a 64-KB allocation unit size.
- **FOLLOW DISK MANAGEMENT BEST PRACTICES,** including stopping the database server instance when removing a data disk or changing the cache type to reduce the risk of database corruption.

AMD EPYC Processors Power Microsoft SQL Virtual Machines in the Cloud

TABLE 2. Da AND Ea SERIES CONFIGURATIONS AND PERFORMANCE CHARACTERISTICS

Da Series	vCPUs	Memory (GB)	Temp Storage (SSD in GB)	Max Data Disks	Max Temp IOPS (MB/sec)	Max Temp Read (MB/sec)	Max Temp Write (MB/sec)	Other
Standard_D2a_v4	2	8	50	4	3000	46	23	Support: • Live migration • Memory preserving updates Do not support: • Premium storage • Premium storage caching
Standard_D4a_v4	4	16	100	8	6000	93	46	
Standard_D8a_v4	8	32	200	16	12,000	187	93	
Standard_D16a_v4	16	64	400	32	24,000	375	187	
Standard_D32a_v4	32	128	800	32	48,000	750	375	
Standard_D48a_v4	48	192	1200	32	96,000	1000	500	
Standard_D64a_v4	64	256	1600	32	96,000	1000	500	
Standard_D96a_v4	96	384	2400	32	96,000	1000	500	

Da Series	vCPUs	Memory (GB)	Temp Storage (SSD in GB)	Max Data Disks	Max Cached and Temp Throughput (IOPS/MBps)	Max Uncached Disk Throughput (IOPS/MBps)	Other
Standard_D2as_v4	2	8	16	4	4000/32 (Cache 50 GB)	3200/48	Support: • Live migration • Memory preserving updates • Premium storage • Premium storage caching
Standard_D4as_v4	4	16	32	8	8000/64 (Cache 100 GB)	6400/96	
Standard_D8as_v4	8	32	64	16	16,000/128 (Cache 200 GB)	12,800/192	
Standard_D16as_v4	16	64	128	32	32,000/255 (Cache 400 GB)	25,600/384	
Standard_D32as_v4	32	128	256	32	64,000/510 (Cache 800 GB)	51,200/768	
Standard_D48as_v4	48	192	384	32	96,000/1020 (Cache 1200 GB)	76,800/1148	
Standard_D64as_v4	64	256	512	32	128,000/1020 (Cache 1600 GB)	80,000/1200	
Standard_D96as_v4	96	384	768	32	192,000/1020 (Cache 2400 GB)	80,000/1200	

Ea Series	vCPUs	Memory (GB)	Temp Storage (SSD in GB)	Max Data Disks	Max Temp IOPS (MB/sec)	Max Temp Read (MB/sec)	Max Temp Write (MB/sec)	Other
Standard_E2a_v4	2	16	32	4	3000	46	23	Support: • Live migration • Memory preserving updates Do not support: • Premium storage • Premium storage caching
Standard_E4a_v4	4	32	64	8	6000	93	46	
Standard_E8a_v4	8	64	128	16	12,000	187	93	
Standard_E16a_v4	16	128	400	32	24,000	375	187	
Standard_E20a_v4	20	160	500	32	30,000	468	234	
Standard_E32a_v4	32	256	800	32	48,000	750	375	
Standard_E48a_v4	48	384	1200	32	96,000	1000	500	
Standard_E64a_v4	64	512	1600	32	96,000	1000	500	
Standard_E96a_v4	96	672	2400	32	96,000	1000	500	

Ea Series	vCPUs	Memory (GB)	Temp Storage (SSD in GB)	Max Data Disks	Max Cached and Temp Throughput (IOPS/MBps)	Max Uncached Disk Throughput (IOPS/MBps)	Other
Standard_E2as_v4	2	16	32	4	4000/32 (Cache 50 GB)	3200/48	Support: • Live migration • Memory preserving updates • Premium storage • Premium storage caching
Standard_E4as_v4	4	32	64	8	8000/64 (Cache 100 GB)	6400/96	
Standard_E8as_v4	8	64	128	16	16,000/128 (Cache 200 GB)	12,800/192	
Standard_E16as_v4	16	128	256	32	32,000/255 (Cache 400 GB)	25,600/384	
Standard_E20as_v4	20	160	320	32	40,000/320 (Cache 500 GB)	32,000/480	
Standard_E32as_v4	32	256	512	32	64,000/510 (Cache 800 GB)	51,200/768	
Standard_E48as_v4	48	384	768	32	96,000/1020 (Cache 1200 GB)	76,800/1148	
Standard_E64as_v4	64	512	1024	32	128,000/1020 (Cache 1600 GB)	80,000/1200	
Standard_E96as_v4	96	672	1344	32	192,000/1020 (Cache 2400 GB)	80,000/1200	

SOLUTION BRIEF
OCTOBER 2020

AMD EPYC Processors Power Microsoft SQL Virtual Machines in the Cloud

CONCLUSION

If you use Microsoft SQL Server in the cloud on Azure Virtual Machines, consider configuring your deployment with AMD EPYC processors. Our Da and Ea Series instances offer the right balance of price/performance to support your move to the cloud. Whether you need to deploy production enterprise workloads or memory-intensive workloads, Azure SQL VMs powered by AMD EPYC processors give you the right combination of vCPUs, memory, and storage with the flexibility to deploy in the cloud without sacrificing performance.

RESOURCES

- AMD AND MICROSOFT AZURE
- AMD EPYC PROCESSORS
- WHAT IS SQL SERVER ON AZURE VIRTUAL MACHINES? *
- DAV4 AND DASV4-SERIES*
- EAV4 AND EASV4-SERIES*
- PERFORMANCE GUIDELINES FOR SQL SERVER IN AZURE VIRTUAL MACHINES*

FOOTNOTES

1. The benchmark workload is derivative and is not directly comparable to published TPC-C results.
2. VM1: Optimized for in-memory hyper-threaded applications E8asv3 (8 vCPU, 64 GiB memory), Local SSD 128 GiB, Premium disk support: Yes, OS: Windows Server® 2020, Premium SSD – 4x (P40, Disk size in GiB: 2048, IOPS per disk : 7500, Throughput per disk: 250 MiB/sec); 2x (P50, Disk size in GiB: 4096, IOPS per disk : 7500, Throughput per disk: 250 MiB/sec), VM2: Optimized for in-memory hyper-threaded applications E16asv3 (32 vCPU, 128 GiB memory), Local SSD 256 GiB, Premium disk support: Yes, OS: Windows Server® 2020, Premium SSD – 8x (P40, Disk size in GiB: 2048, IOPS per disk : 7500, Throughput per disk: 250 MiB/sec); 2x (P50, Disk size in GiB: 4096, IOPS per disk : 7500, Throughput per disk: 250 MiB/sec), VM3: Optimized for in-memory hyper-threaded applications E32asv3 (32 vCPU, 256 GiB memory), Local SSD 512 GiB, Premium disk support: Yes, OS: Windows Server® 2016, Premium SSD – 8x (P40, Disk size in GiB: 2048, IOPS per disk : 7500, Throughput per disk: 250 MiB/sec); 2x (P50, Disk size in GiB: 4096, IOPS per disk : 7500, Throughput per disk: 250 MiB/sec), Performance Measured with SQL Server® 2020, HammerDB-3.1.
3. Microsoft: Performance Guidelines for SQL Server on Azure Virtual Machines*, October 18, 2019.

* Links to third party sites are provided for convenience and unless explicitly stated, AMD is not responsible for the contents of such linked sites and no endorsement is implied.

LE-73401-00 10/20