

# Selecting the Right Size Azure VM

Tim Radney  
tim@timradney.com

# Speaker: Tim Radney



- Consultant/Trainer/Speaker/Author
- SaaS Migration Lead / Principal Consultant
- Email: [Tim@SQLskills.com](mailto:Tim@SQLskills.com) and [Tim@timradney.com](mailto:Tim@timradney.com)
- Blog: <https://www.SQLskills.com/blogs/Tim>
- Blog: <http://www.timradney.com>
- Microsoft Data Platform MVP
- Chapter Leader "Cloud Data Platform Virtual Group"

## Key technology areas:

- Azure SQL Virtual Machine and storage
- SQL Server performance, tuning and optimization
- Azure Data Services Security
- Disaster Recovery
- Azure SQL DB / Managed Instance



 @TRADNEY



# Azure SQL Virtual Machines

- What is Microsoft Azure
- Introduction to IaaS
- Getting started with Azure VMs
- High availability and disaster recovery
- Azure VM size considerations
- Capturing key performance metrics



# Microsoft Azure

A cloud computing platform and infrastructure created by Microsoft, for building, deploying and managing applications and services through a global network of Microsoft-managed and Microsoft partner hosted datacenters.

[https://en.wikipedia.org/wiki/Microsoft\\_Azure](https://en.wikipedia.org/wiki/Microsoft_Azure)

# IaaS



- Infrastructure as a Service
  - Networking
  - Security
  - Compliance
  - Storage
  - Virtual Machines



# IaaS Networking

- Point-to-Site VPN – (P2S VPN)
  - Single computer (point) to Azure network (site)
  - Uses certificates
  - Adds flexibility for roaming users
  - Consider Active Directory Certificate Services (AD CS)
- Site-to-Site VPN – (S2S VPN)
  - Persistent connection from on-prem to Azure network
  - Requires a static public IP address
- ExpressRoute
  - Fastest connection 50Mbps to 10Gbps
  - Limited availability



# Virtual Private Networking

- Internal networks
  - VNets – Virtual Networks
    - Can create multiple VNets to isolate environments
    - Can subdivide using Subnets
    - Setup VPNs to allow VNets to communicate
- Design and plan first
  - Before creating VMs or other services, setup your VNet



# Security

- Network Security Groups (NSGs)
  - Enforce and control network traffic rules
  - Similar to a firewall
- Azure Active Directory
  - Manage identity and access
  - Integrated for both cloud and on-premises




# Compliance

- Many certifications
  - DoD Provisional Authorizations at Impact Levels 5, 4, and 2
  - FIPS 140-2 – US Federal Info Processing Standards
  - HIPPA/HITECH – Health Care
  - ISO 22301, 27001, 27017, 27018
  - PCI DSS – Payment Card industry
  - CJIS – US Criminal Justice Information Services
  - EU Model Clauses
  - And more..
- Microsoft compliance offerings: <http://bit.ly/2lKjwzK>




# Storage




- VM starts with C: and D:
  - C: registered as SATA, labeled /dev/sda by default
  - Max capacity of C: limited to 1TB
  - D: is temporary, nothing important should be placed on it
  - D: is ideal for page or swap files and tempdb. Size varies on size of VM
- VM size matters
  - VM size determines number of disk, disk types, throughput, and more
    - Standard – HDD
    - Standard - SSDs
    - Premium – SSDs
    - Ultra disk

# Storage



- Unmanaged disks
  - Traditional disk
  - You manage
  - Don't allocate too many disks on the same account
- Managed disks
  - Creation and management handled for you
  - Specify size and performance tier (Standard/Premium)
- Scale
  - Add more disks and stripes (RAID)

# Storage



- Standard Disk

- HDD
- Cost effective for dev/test
- 4GB – 32TB
- 60 – 750 MB/s throughput and 500 – 6,000 IOPS per disk


- Premium Disk

- SSD – high performance
- 4GB – 32TB
- 25 MB/s – 900MB/s throughput
- 120 – 20,000 IOPS per disk


# Virtual Machines

- Creating the VM

- Huge inventory of images
  - OS based, solution templates, application infrastructure, and business applications
- Image with OS and SQL Server installed
  - Licensed or BYOL (Bring Your Own License)
  - Everything is installed
  - Some time spent cleaning, uninstalling
- Image with OS, you install and license SQL Server
  - You specify what gets installed and where
  - Higher chance of not configuring storage properly




# Sizing Virtual Machines



- You get predictable performance
- Memory, vCore, and throughput options you can scale
- Sizes that support higher memory to core ratios
- Various VM families for specific VM performance needs
- SQL Server typically fits within Memory optimized, General Purpose, and Storage optimized

# Naming convention explanation



- **Family** and Sub-family - Indicates the VM Family Series and specialized variations
- # of **vCPUs** - Denotes the number of vCPUs of the VM
- **Features:**
  - a = AMD-based processor
  - b = bandwidth (highest I/O throughput available in Azure VMs)
  - d = diskful (local temp disk is present); this is for newer Azure VMs, see Ddv4 and Ddsv4-series
  - i = isolated size
  - l = low memory; a lower amount of memory than the memory intensive size
  - m = memory intensive; the most amount of memory in a particular size
  - t = tiny memory; the smallest amount of memory in a particular size
  - s = Premium Storage capable, including possible use of Ultra SSD (Note: some newer sizes without the attribute of s can still support Premium Storage e.g. M128, M64, etc.)
- Accelerator Type - Denotes the type of hardware accelerator in the specialized/GPU SKUs.
- Version - Denotes the version of the VM Family Series

# Naming convention example

[Family] + [Sub-family]\* + [# of vCPUs] + [Additive Features] + [Accelerator Type]\* + [Version]

## Virtual Machine: E64-16ds\_v4

Value	Details
Family	E
# of vCPUs	64
Available Cores	16
Additive Features	d = temporary / ephemeral storage s = Premium Storage capable

VM Naming Convention


<https://docs.microsoft.com/en-us/azure/virtual-machines/vm-naming-conventions>

Constrained vCPU VMs

<https://docs.microsoft.com/en-us/azure/virtual-machines/windows/constrained-vcpu>




# High Availability and Disaster Recovery




- HA/DR options
  - Always On Availability Groups
  - Always On Failover Cluster Instances
  - Log shipping
  - Database mirroring
  - SQL Server backup and restore with Azure Blob storage service

# High Availability and Disaster Recovery



- Azure-only Availability Groups
  - Currently requires a WSFC
  - Due to WSFC, must have a domain controller
  - Servers will need to be in the same resource group
  - Select your service name carefully
- Hybrid
  - Easy to extend on-premises AG with a secondary in Azure
  - Requires VPN, extend AD to Azure
  - Inexpensive DR solution


# High Availability and Disaster Recovery



- Failover Cluster Instances


- Requires using Windows Server 2016 Storage Spaces Direct (S2D)
- Third party tool – SIOS Datakeeper
- Remote iSCSI Target shared blocks storage via ExpressRoute
- Azure File storage

# High Availability and Disaster Recovery




- Log shipping
  - Just like on-premises
  - Primary or secondary is in Azure
  - Requires VPN
  - Allows for load delay
- Database mirroring
  - Cross-site disaster recovery using server certificates
    - Doesn't require same domain or VPN
  - Cross-site with Active Directory would require VPN
  - Deprecated in 2016

# High Availability and Disaster Recovery




- SQL Server backup with Azure Blob Storage Service
  - Backup on-premises to Azure Blob Storage
  - 2016 enhancements
    - Striping to multiple blobs for support up to 12.8TB
    - Snapshot backup – nearly instantaneous backups and rapid restores for db files
    - Managed Backup Scheduling – custom schedules

# High Availability and Disaster Recovery



- HA/DR options
  - Always On Availability Groups
  - Always On Failover Cluster Instances
  - Log shipping
  - Database mirroring
  - SQL Server backup and restore with Azure Blob storage service

# Azure VM Size Considerations



- Understand your workload
  - CPU utilization
  - Memory utilization
  - Storage space requirement
  - Disk usage metrics
    - I/O
    - Throughput\*
- Size for what you need and use
  - You're essentially renting, scale when you grow

# Recommended SQL VM Series

## Edsv5-series

- Best all-purpose SQL Server virtual machine
- Good remote throughput
- Large local drives with high local/cached throughput
- **Offers best price-performance for SQL Server workloads**

## M-series\*

- Highest I/O throughput and memory available in Azure
- Great for mission critical OLTP and high-end data warehouse workloads
- Expensive – it is recommended to start with Edsv4 first if possible

## Ddsv5-series

- Entry level SQL Server virtual machine
- Good choice for dev/test workloads
- Works for smaller production database environments



# Recommended SQL VM Series

Edsv5-series

- Best all-purpose SQL Server virtual machine
- Good remote throughput
- Large local drives with high local/cached throughput
- **Offers best price-performance for SQL Server workloads**

When in doubt, pick this one!

# Recommended SQL VM Memory Optimized Series


Ebdsv5-series

- Highest remote I/O throughput available in Azure
- Large local drives with high local/cached throughput
- **Offers the best price-performance for SQL Server workloads**

Or better yet, this one which went GA yesterday!

# Azure VM Size Considerations

- Azure VMs have limits on
  - Number of cores
  - Amount of memory
  - Ephemeral disk support and size
  - Number of disk
  - I/O and throughput
  - Network bandwidth



# VM I/O Limits

Size	vCPU	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
Standard_E2ds_v4	2	16	75	4	19000/120(50)	3200/48
Standard_E4ds_v4	4	32	150	8	38500/242(100)	6400/96
Standard_E8ds_v4	8	64	300	16	77000/485(200)	12800/192
Standard_E16ds_v4	16	128	600	32	154000/968(400)	25600/384
Standard_E20ds_v4	20	160	750	32	193000/1211(500)	32000/480
Standard_E32ds_v4	32	256	1200	32	308000/1936(800)	51200/768
Standard_E48ds_v4	48	384	1800	32	462000/2904(1200)	76800/1152
Standard_E64ds_v4	64	504	2400	32	615000/3872(1600)	80000/1200

<https://docs.microsoft.com/azure/virtual-machines/edv4-edsv4-series>

# New Ebdsv5 VM Series

Higher I/O throughput at lower core counts, purposefully designed with RDBMS in mind

Generally available today!!!!

~~Improves price-perf by 30% compared to Edsv4~~  
Improves price-perf by 34% compared to Edsv4

Size	vCPU	Memory: GiB	Max uncached disk throughput: IOPS/MBps
Standard_E2bds_v5	2	16	5500/156
Standard_E4bds_v5	4	32	11000/350
Standard_E8bds_v5	8	64	22000/625
Standard_E16bds_v5	16	128	44000/1250
Standard_E32bds_v5	32	256	88000/2500
Standard_E48bds_v5	48	384	120000/4000
Standard_E64bds_v5	64	512	120000/4000

TPC-C	E64-32ds_v4	E32bds_v5
<b>Storage Config</b>	Data pool, 12 P30s Log pool, 4 P30s	Data pool, 18 P30s Log pool, 2 P30s
<b>Compute Cost*</b>	\$3,784.32/month	\$1,950.56/month
<b>Storage Cost</b>	\$2,162.72/month	\$2,703.40/month
<b>Total Cost</b>	\$5,947.04/month	\$4,653.96/month
<b>NOPM</b>	1,065,949	1,259,418
<b>Price/Perf Ratio*</b>	<b>\$5.57</b> /1000 NOPM	<b>\$3.70</b> /1000 NOPM

\* This cost is for pay-as-you-go compute only, assuming Azure Hybrid Benefit for both Windows OS and SQL Server licensing costs.

<https://aka.ms/AzureSQLVMPricePerf>

# Optimize SQL Server license cost with constrained vCPU VMs




Name	vCPU	Specs
Standard_M8-2ms	2	Same as M8ms
Standard_M8-4ms	4	Same as M8ms
Standard_M16-4ms	4	Same as M16ms
Standard_M16-8ms	8	Same as M16ms
Standard_M32-8ms	8	Same as M32ms
Standard_M32-16ms	16	Same as M32ms
Standard_M64-32ms	32	Same as M64ms
Standard_M64-16ms	16	Same as M64ms
Standard_M128-64ms	64	Same as M128ms
Standard_M128-32ms	32	Same as M128ms

Azure offers VM sizes where you can constrain the vCPU count to reduce the cost of software licensing, while maintaining the same **memory**, **storage**, and **I/O bandwidth**


<https://docs.microsoft.com/en-us/azure/virtual-machines/windows/constrained-vcpu>

# VM Size – SQL Server

- Size determines IOPS and throughput
  - E2s\_v3 – 2 vCPU, 16GB RAM, 3200 IOPS, 48 MBps
  - E4-2s\_v3 – 2 vCPU, 32GB RAM, 6400 IOPS, 96 MBps
  - E8-2s\_v3 – 2 vCPU, 64GB RAM, 12,800 IOPS, 192 MBps
- Benchmark using an 8TB SSD with 16,000 IOPS and 500 MBps
  - Smaller sized VMs will restrict IOPS and throughput regardless of the capabilities of the storage provisioned!



# P50 disk limiting to 250MB/sec



Task Manager Performance - Disk 4 (Z:) Microsoft Virtual Disk

Disk transfer rate: 250 MB/s

Active time: 100%

Read speed: 254 MB/s | Write speed: 1.1 MB/s

Capacity: 3.0 TB | Formatted: 3.0 TB

System disk: No | Page file: No

Active time: 100% | Average response time: 1005 ms

2013.6348 (Disk may be limiting speed)

2047.7593 (Disk may be limiting speed)

2024.6762 (Disk may be limiting speed)

2034.0073 (Disk may be limiting speed)

2041.9038 (Disk may be limiting speed)

2049.4363 (Disk may be limiting speed)

2040.8528 (Disk may be limiting speed)

1992.5268 (Disk may be limiting speed)

2091.4535 (Disk may be limiting speed)

2020.4152 (Disk may be limiting speed)

1996.5087 (Disk may be limiting speed)

2085.3219 (Disk may be limiting speed)

2017.5585 (Disk may be limiting speed)

2027.2708 (Disk may be limiting speed)

2060.4877 (Disk may be limiting speed)

2039.9208 (Disk may be limiting speed)

2016.8039 (Disk may be limiting speed)

2063.1257 (Disk may be limiting speed)

2035.3014 (Disk may be limiting speed)

1999.3354 (Disk may be limiting speed)

2060.9992 (Disk may be limiting speed)

2047.1078 (Disk may be limiting speed)

2013.4932 (Disk may be limiting speed)

2036.2363 (Disk may be limiting speed)

2026.2907 (Disk may be limiting speed)

2051.415 (Disk may be limiting speed)

2071.1175 (Disk may be limiting speed)

2000.2866 (Disk may be limiting speed)

2076.0611 (Disk may be limiting speed)

2008.3736 (Disk may be limiting speed)

2049.1663 (Disk may be limiting speed)

1993.7989 (Disk may be limiting speed)

2051.5627 (Disk may be limiting speed)

2052.7272 (Disk may be limiting speed)

2042.3784 (Disk may be limiting speed)

2041.8685 (Disk may be limiting speed)

2023.7975 (Disk may be limiting speed)

2050.9252 (Disk may be limiting speed)

2033.6943 (Disk may be limiting speed)

2087.0009 (Disk may be limiting speed)

1998.7652 (Disk may be limiting speed)

2020.6454 (Disk may be limiting speed)

2044.2262 (Disk may be limiting speed)

1999.5355 (Disk may be limiting speed)

2104.7638 (Disk may be limiting speed)

1997.5745 (Disk may be limiting speed)

2038.2676 (Disk may be limiting speed)

2054.6678 (Disk may be limiting speed)

2040.6597 (Disk may be limiting speed)

9.5 %	2554	Done	0	Failed	4122	Pending	0	Skipped	6676	Total	2-sec	Throughput	(Mb/s)	2033.6943	(Disk may be limiting speed)
9.5 %	2558	Done	0	Failed	4118	Pending	0	Skipped	6676	Total	2-sec	Throughput	(Mb/s)	2087.0009	(Disk may be limiting speed)
9.5 %	2564	Done	0	Failed	4112	Pending	0	Skipped	6676	Total	2-sec	Throughput	(Mb/s)	1998.7652	(Disk may be limiting speed)
9.5 %	2568	Done	0	Failed	4108	Pending	0	Skipped	6676	Total	2-sec	Throughput	(Mb/s)	2020.6454	(Disk may be limiting speed)
9.6 %	2576	Done	0	Failed	4100	Pending	0	Skipped	6676	Total	2-sec	Throughput	(Mb/s)	2044.2262	(Disk may be limiting speed)
9.6 %	2582	Done	0	Failed	4094	Pending	0	Skipped	6676	Total	2-sec	Throughput	(Mb/s)	1999.5355	(Disk may be limiting speed)
9.6 %	2589	Done	0	Failed	4087	Pending	0	Skipped	6676	Total	2-sec	Throughput	(Mb/s)	2104.7638	(Disk may be limiting speed)
9.6 %	2597	Done	0	Failed	4079	Pending	0	Skipped	6676	Total	2-sec	Throughput	(Mb/s)	1997.5745	(Disk may be limiting speed)
9.6 %	2603	Done	0	Failed	4073	Pending	0	Skipped	6676	Total	2-sec	Throughput	(Mb/s)	2038.2676	(Disk may be limiting speed)
9.7 %	2609	Done	0	Failed	4067	Pending	0	Skipped	6676	Total	2-sec	Throughput	(Mb/s)	2054.6678	(Disk may be limiting speed)
9.7 %	2613	Done	0	Failed	4063	Pending	0	Skipped	6676	Total	2-sec	Throughput	(Mb/s)	2040.6597	(Disk may be limiting speed)



# Monitoring I/O Capping

## Metrics that help diagnose disk I/O capping

- Data Disk IOPS Consumed Percentage
- Data Disk Bandwidth Consumed Percentage

## Metrics that help diagnose VM I/O capping

- VM Cached IOPS Consumed Percentage
- VM Cached Bandwidth Consumed Percentage
- VM uncached IOPS Consumed Percentage
- VM Uncached Bandwidth Consumed Percentage

# Azure SQL Virtual Machine storage options

## 1. Choose the right disk

Premium Disk



P30s, P40s for data, P30+ for log files

Ultra Disk



Ultra Disk for low disk latencies

Ephemeral Disk



Use ephemeral disk for tempdb

## 2. Enable host cache

Data files



Read-only

Log files



None

tempdb



Read-only when remote

## 3. Enable Write Acceleration

M-series only



Enabled like caching in the Azure portal

Beneficial for Log files



Improved disk latencies

Disk support varies



Refer to M-series documentation

## 4. Leverage Storage Configuration (Best Practices)

Azure Portal



Enabled in the Azure portal

Built-in Best Practices




Blocksize, interleave, numberofcolumns

Disk location & Features



File separation, network acceleration

# Key Metrics to Consider



- CPU
  - # of vCore
  - Consider source speed to VM
  - Are you over/under provisioned
  - Can you tune to better optimize
- Memory
  - How much is enough
  - Are you over/under provisioned
  - Can you tune to better optimize
- Disk
  - Storage capacity
  - I/O needs – this alone is over 50% of performance support cases!

# Measure Application Performance Requirements

Counter	Description	PerfMon (Windows)	lstat (Linux)
★ ★ ★ IOPS or Transactions per second	Number of I/O requests issued to the storage disk per second.	Disk Reads/sec	tps
		Disk Writes/sec	r/s
			w/s
★ ★ ★ Throughput	Amount of data read from or written to the disk per second.	Disk Read Bytes/sec	kB_read/s
		Disk Write Bytes/sec	kB_wrtn/s
★ ★ ★ Latency	Total time to complete a disk IO request.	Average Disk sec/Read	await
		Average disk sec/Write	svctm


Uncached IOPS/MBps

- Roughly maps to transaction log reads and writes + data file writes and ~ 70% of reads (depends on host cache hit ratio)

Cached and local IOPS/MBps

- Roughly maps to tempdb reads and writes + data file reads and writes

# Resources




- SQLDiag: <https://docs.microsoft.com/en-us/sql/tools/sqldiag-utility>
- SQLNexus: <https://github.com/Microsoft/SqlNexus>
- Perfmon: <https://docs.microsoft.com/en-us/windows-server/administration/windows-commands/perfmon>
- PAL: <https://github.com/clinthuffman/PAL>
- Perfinsights: <https://docs.microsoft.com/en-us/troubleshoot/azure/virtual-machines/performance-diagnostics>
- SQLinsights: <https://docs.microsoft.com/en-us/azure/azure-monitor/insights/sql-insights-overview>

# Demo:

---

Azure Portal – step through provisioning an Azure VM

A perspective view of a server room aisle. The racks on both sides are filled with server units, many of which have small, glowing lights (green and white) on their front panels. The floor is a light-colored, perforated metal grating. The ceiling has recessed lighting and visible ductwork. The overall atmosphere is clean, modern, and high-tech.

**Thank You!**  
**@tradney**  
**tim@timradney.com**