

SQLskills Immersion Event

IEAzure: Azure VMs and Azure SQL Database

Module 3: Migrating to Azure Virtual Machines

Tim Radney

Tim@SQLskills.com



Overview

- Collecting on-premises data for sizing
- Benchmarking Azure virtual machines
- Choosing the correct size virtual machine
- Migration strategies
- Pros and cons

Collecting On-Premises Data for Sizing

- **What size virtual machine are you going to need**
- **You need to benchmark your current environment first**
 - CPU
 - Memory
 - Disk IO
 - Storage requirements
- **Physical to virtual concerns**
 - Many physical machines are overpowered by design, you purchase for room to scale up
 - Replacing physical machines are costly, most organizations over scale to support 3-5 years of growth
 - VMs are easier to scale, size more for 12 months to save on operational cost and license

Collecting On-Premises Data for Sizing

- **Virtual machines are already categorized and optimized for CPU, memory, storage.**
 - General purpose – balanced CPU-to-memory ratio. Ideal for testing and development, small to medium databases, and low to medium traffic web servers
 - B, Dsv3, Dv3, DSv2, Dv2, Av2, DC
 - Memory optimized – high memory-to-CPU ratio. Great for relational database servers, medium to large caches, and in-memory analytics.
 - Esv3, Ev3, M, GS, G, DSv2, Dv2
 - Storage optimized – high disk throughput and IO ideal for Big Data, SQL, NoSQL databases, data warehousing and large transactional databases.
 - Lsv2 and Ls
- **Most SQL Servers fall into DSv2 or Dv2**

CPU

- **How much CPU do you need?**
 - It depends on current workload
- **You need to capture current CPU workload**
 - Third-party tools
 - DMV scripts
 - Glenn Berry diagnostic queries - <http://bit.ly/2rXT0ar>
 - CPU utilization history for last 256 minutes
 - Look for average CPU utilization during your peak usage times
- **Capture current CPU specifications and benchmark**
 - CPU-Z
- **Considerations before capturing data**
 - Power plan – are you running in High Performance?

Memory

- **How much memory (ram) do you need?**
 - It depends on workload and performance requirements
- **You need to capture current memory usage patterns**
 - Page Life Expectancy, Available Mbytes, Max Server Memory configuration
 - Third party tools
 - DMVs, Perfmon, queries

Storage

- **What are your current IO patterns?**
 - Workloads dictate this
- **Capture current usage patterns**
 - Disk latency – how long is it waiting?
 - Reads/Writes – throughput
 - Current benchmarks
 - CrystalDiskMark
 - Microsoft DiskSpd
- **Storage can easily be a bottleneck**
 - Indexing
 - Faster IO
 - More memory (bufferpool)
- **You may need to tune before migrating**

General Concerns

- **Physical-to-virtual and virtual-to-virtual concerns**
 - It is easy to over allocate memory on physical machines
 - Memory is somewhat inexpensive
 - Fairly easy to get lots of memory on on-premises VMs
 - Memory is a valuable resource in the Cloud
 - Higher memory = higher core count = more SQL Server license

Data Collection Tools

■ Free methods

- Perfmon counters
- DMVs
 - <http://bit.ly/2rXT0ar>
- Tigertoolbox – SQL Performance Baseline
 - <https://bit.ly/2YEKWAd>

■ Paid tools

- SentryOne
 - <https://www.sentryone.com/plan-explorer>

■ Importance of baselines

- <https://sqlperformance.com/2018/05/baselines/importance-of-baselines>

Demo

Capturing CPU, memory, and disk benchmarks

Benchmarking Azure Virtual Machines

- **Is collecting baseline/benchmark data different in Azure virtual machines than on-premises?**
 - No, not really
 - You don't have VMware/Hyper-V console access
 - Everything else is mostly the same
- **Capture benchmarks – what the system is capable of**
 - CPU-Z
 - CrystalDiskMark – Microsoft DiskSpd
 - Create a virtual machine to your specifications, run your benchmarks, shutdown and deallocate the virtual machine until your next test or destroy

Demo

Running benchmarks against Azure VMs

Choosing the Correct Size VM

- **You decide on VM by vCPU, memory, disk**
 - vCPU and memory are easy to scale by migrating to a larger VM
 - Larger VM sizes allow for a higher number of data disk
 - More data disk allow you to stripe more disk for more IOPS
 - Modifying disk structures for more IO is much more complex than just moving to a larger VM
- **Lower core count with higher ram sizes exist with memory optimized models**
 - 4 vCPU and 128GB, 8 vCPU and 256GB, 16 vCPU and 432GB are your lowest vCPU and max memory available
- **Pay attention to throughput limits on VM size**
 - It is entirely possible to provision storage with higher IOPS and throughput than the VM limits allow

Demo

Azure Marketplace: sizing various VMs looking at vCPU combinations with memory and data disk, VM limits

Microsoft Database Experimentation Assistant

- DEA is an A/B testing solution for SQL Server upgrades
- Evaluates a version of SQL for a given workload
- Works for SQL Server 2005 and above
- Provides an automated setup for a workload capture and replay of a production database. This uses Distributed Replay and SQL Server-side tracing
- Provides a rich user experience by visualizing data through an analysis report
- Data can be exported to a *.csv file for sharing
- Supports Windows 7, Windows 8, Windows 8.1, Windows 10, Windows Server 2012/2016/2019

Demo

Installing, configuring, and using DEA

Migration Strategies

- **SQL Server backup with Azure Blob storage service**
 - Backup to URL
 - Backup files are available to the VM for easy restore
 - Good for staging a full restore
 - DR
- **Backup and restore**
 - Traditional method for migration
 - Backup local and copy to Azure Blob storage, restore
- **Log shipping**
 - Setup VPN or connectivity to Azure VM
 - Establish log shipping
 - Quick cutover
 - HA/DR

Migration Strategies

- **Database mirroring**

- Setup VPN or connectivity to VM
- Setup mirroring like you normally do
- Easy failover
- HA/DR

- **Windows Import/Export Service – ship a hard drive**

- Used for large migrations
- Carrier transfer is faster than data speeds available to some organizations
- Requires planning to sync data changes

- **Availability Group Replica**

- Extend an on-premises availability group to Azure
- Extend AD to Azure AD, setup VPN
- Used for DR for numerous organizations (Hybrid)

Migration Strategies

- **Database Migration Assistant**
 - UI to assist with backup and restore
- **Detach, copy to Azure blob, attach**
 - Traditional and risky method to detach a database, copy files, and attach
 - This method is not recommended
 - Backup and restore is much safer
- **Physical-to-virtual (P2V), VHD to VHD**
 - Tools exist make an image of a physical machine
 - Migrate an existing VM by migrating the VHD to Azure
- **Transactional replication**
 - Setup a subscriber on an Azure VM, fail over

Demo

Migration strategies from on-premises to an Azure VM

Pros of Azure VMs

- Affordable
- Infrastructure is handled for you
- Compliant
- Secure
- Integrates well
- Hybrid is an option

Pros of Azure VMs

- **Geo-redundant datacenters**
- **Give up some of your control of the systems**
 - Infrastructure roles, servers, storage, network
- **High SLAs**
 - 99.9% uptime - single instance VM
 - 99.95% uptime - VMs with two or more instances in the same Availability Set
 - 99.99% uptime - VMs with two or more instances in two or more Availability Zones in the same Azure region

Cons

- **Cost can seem high for larger compute models**
 - Total cost of ownership is likely lower
- **Give up some of your control of the systems**
 - Infrastructure roles, servers, storage, network
- **Latency with hybrid solutions**
- **Fear**

Key Takeaways

- Working with Azure virtual machines isn't much different than any other virtual machine
- Migrating to Azure VMs is the same process as any other physical-to-virtual migration or VM to VM migration
- There are numerous ways of migrating from on-premises to Azure VMs

Review

- How to collect on-premises data for sizing
- What benchmarks and how to benchmark Azure virtual machines
- How to choose the correct size virtual machine
- Numerous migration strategies
- Different pros and cons of Azure virtual machines

References

- **CPU-Z**
 - <http://cpuid.com>
- **CrystalDiskMark**
 - <http://bit.ly/2vfKjJp>
- **Microsoft DiskSpd**
 - <http://bit.ly/1NkG1BM>
- **Glenn Berry's DMV queries**
 - <http://bit.ly/2rXT0ar>

Questions?

