

Azure VM Performance Considerations Running SQL Server

Vinod Kumar M

@vinodk_sql

<http://blogs.ExtremeExperts.com>

*Your
company
logo here*

saltmarch
MEDIA

GREAT INDIAN
DEVELOPER
SUMMIT



Session Objectives And Takeaways

Session Objective(s):

Learn the performance characteristics of running SQL Server on Windows Azure VM (IaaS)

Understand the performance best practices for running SQL Server on Windows Azure VM

Key Takeaway 1

What impacts performance of SQL Server running on Windows Azure VM

Key Takeaway 2

How to monitor and troubleshoot performance of SQL Server running on Windows Azure VM

Performance tuning

- Analyzing performance impact
 - Is it a problem with my SQL application?or
 - Is it a problem with how I configured the Azure IaaS platform for my usage

Key SQL Server Performance Considerations

Windows Azure IaaS Characteristics

Best Practices

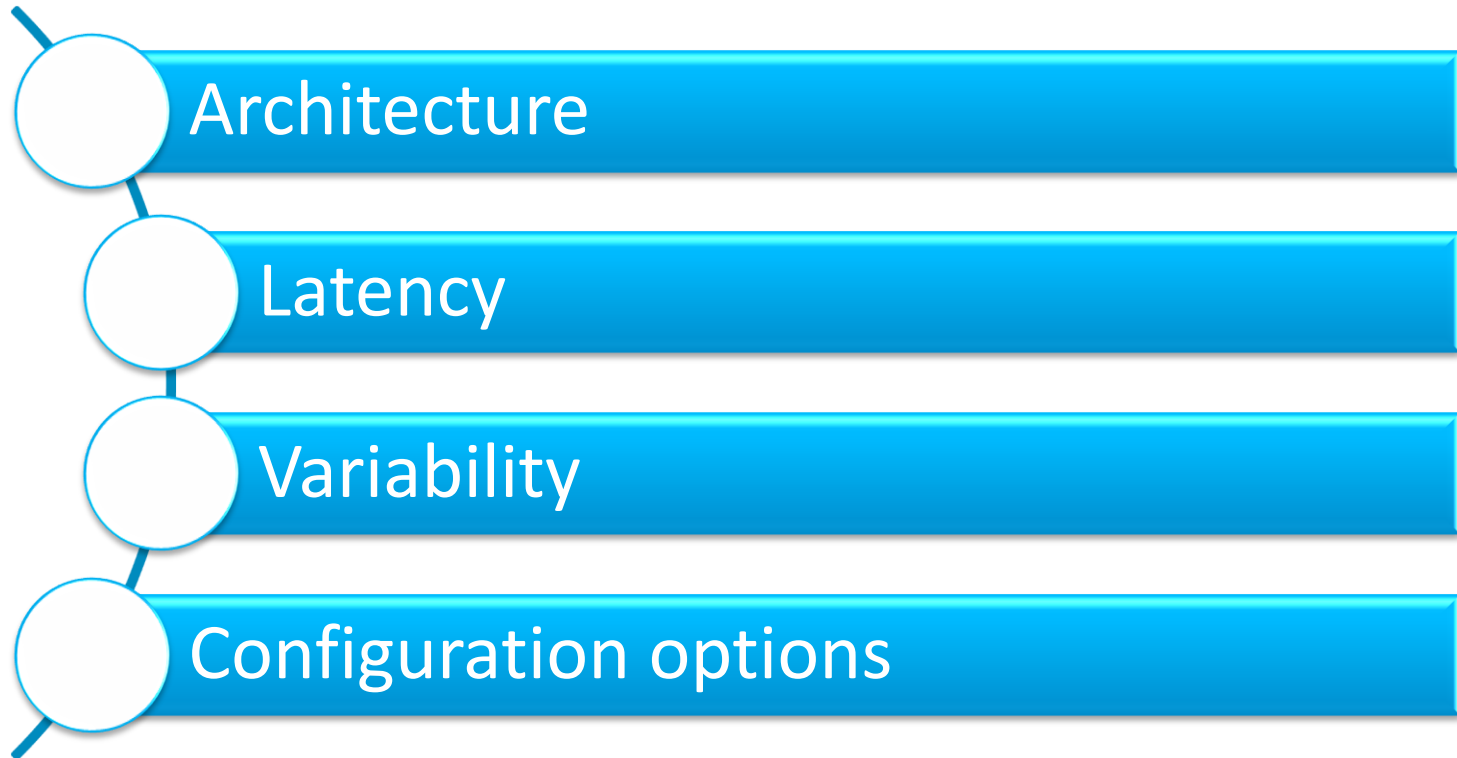
Monitoring

Key SQL Server Performance Considerations


- KPIs
 - Throughput
 - Response time (aka latency)

Metric	OLTP	DW	Transaction Log
Read/Write mix	<ul style="list-style-type: none">• Large number of small transactions• Transactions are mostly similar in pattern• Significant amount of reads	<ul style="list-style-type: none">• Scan intensive, large portions of data at a time• Bulk loading	<ul style="list-style-type: none">• Mostly writes• Requires low latency
IO size and pattern	<ul style="list-style-type: none">• Between 8 and 64K• Mostly random	<ul style="list-style-type: none">• 1 64KB read per 8 512KB reads• Mostly sequential• MB/s a critical metric	<ul style="list-style-type: none">• Highly sequential
# users	high	low	n/a

Windows Azure VM Characteristics




VM Configuration Options




Virtual Machine
size



Network
bandwidth



Disk types and
configurations



Disk cache
settings

VM Disk Types & Configurations

- OS disk (persistent)
 - 127GB disk optimized for OS access patterns (e.g. boot up)
- Data disk (persistent)
 - A VHD you can attach to a VM to store app data
 - Up to 1TB in size
 - Up to 16 disks for XL VMs
- Temporary local disk (non-persistent)
 - Used for transient/temporary data storage & OS page files
 - Hosted in attached disks on physical host
 - Cleaned up in the event of a VM failure or recycling
 - Physical disks shared across other VMs on same physical machine
 - Not recommended for user or system database files

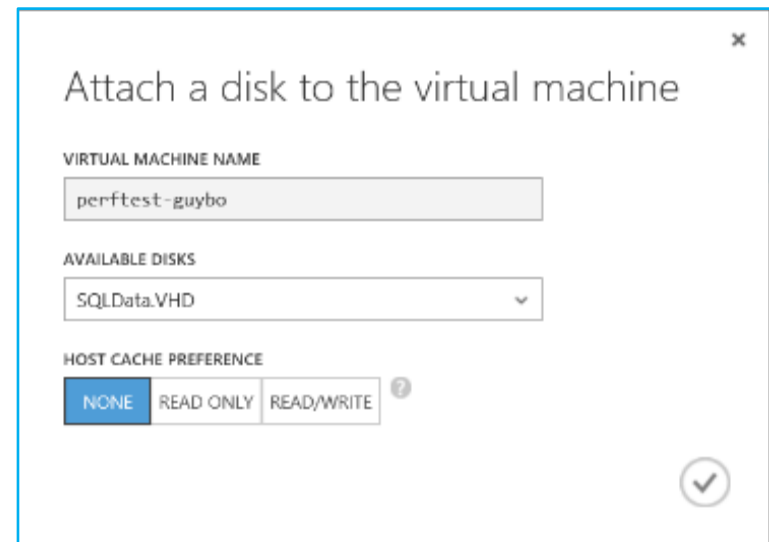
What are Disk Caching Best Practices for SQL Server?

OS Disk

- “Read Write” (default) reduces read latency for IO intensive workloads with smaller DBs (<=10GB)
 - Working set can fit in disk cache or memory, reducing blob storage IO

Data disks

- Cache setting depends on the IO pattern and workload intensity
- Use default of “None” (disable) for higher rate of random IOs (e.g. OLTP) & higher throughput
 - Bypasses physical host local disks, maximizing IO rate
- Consider enabling read cache for latency sensitive read heavy workloads



Best Practices



Network
Latency



Disk
configuration



TempDB
placement



IO
optimization



Should TEMPDB go on D: drive?

- Short answer: No
- Long answer: No
- Why?
- Predictable performance:
OS or data disk can provide same or better performance but D: drive can be more variable, being a physical disk sharing IO with other VMs on the host. Size and performance also depends on VM size
- Configuration overhead:
SQL Server has to recreate TEMPDB in D: if VM goes down – SQL Server service account requires Admin privileges. If stored in a separate folder this needs to be created at startup.
- TEMPDB can be critical to application performance
 - Follow [tempDB IO best practices](#)

Instant File Initialization - Recommendations

- Pre-size all database files appropriately
 - Autogrow is insurance
 - Not a license for drunk driving
- Configure Instant File Initialization before creating database
- Configure Instant File Initialization before extending files.
Example: TEMPDB
- Configure Instant File Initialization before restoring database on a new server
- Don't forget to restart SQL Server after configuring Instant File Initialization

What else to consider for data disk performance?

- Disk warm-up
- NTFS Allocation Unit Size ?
- Single vs. multiple storage accounts with a single VM
 - DO NOT SPREAD DATA FILES OF A SINGLE DATABASE INTO MULTIPLE STORAGE ACCOUNTS !!!
 - Data in different blobs not written at the same time
 - BLOBs that make up the stripe set could be out of sync
 - Instead:
 - Spread the data files across multiple disks to achieve higher IOPS / bandwidth

Note: a storage account has a limit of 20K tps

Performance Monitoring & Troubleshooting

Key tools

Storage
Analytics

VM
Dashboard

SQL Perf
counters

Conclusions

- Read the [Performance Guidance for SQL Server in Windows Azure Virtual Machines](#) white paper
- Plan and test for IO perf variability
- Identify optimal VM size
- Optimize for reduced IO and network round trips
- Use Filegroups and multiple data disks for large DBs
- Identify your KPIs to monitor
- Revisit optimization decisions as workload changes

Thank you!

saltmarsh

GREAT INDIAN
DEVELOPER
SUMMIT

