

Microsoft SQL Server 2016 on Hyper Converged Infrastructure using VMware vSphere 6

TECHNICAL WHITE PAPER



Contents

Executive summary	3
Introduction	3
HPE Hyper Converged 380 family	4
Solution overview	4
Solution components	6
Hardware	6
Software	8
Application software	8
Best practices and configuration guidance for the solution	8
Capacity and sizing	10
Scale-up scenario	
Scale-out scenario	
Consolidation scenario	14
Analysis and recommendations	18
Summary	19
Resources and additional links	19



Executive summary

Many IT organizations are standardized on a "Virtualize First" policy with VMware® vSphere, which drives IT architects to look for configurable, scalable, agile, and highly available hyper converged systems. HPE Hyper Converged 380 combines Hewlett Packard Enterprise's leading server technologies with a software stack that includes HPE StoreVirtual VSA software-defined storage and the HC 380 Management UI and automation software. The result is a solution that combines enterprise-class features and scalability with an easy-to-use software stack for virtualizing your business-critical Microsoft® SQL Server for mid-size enterprises.

The solution described in this paper provides customers with the expected scale-up and scale-out performance implications and database consolidation associated with deploying Microsoft SQL Server OLTP workloads on VMware vSphere 6 using HPE Hyper Converged 380.

Testing shows the HPE Hyper Converged 380 is able to service a range of Microsoft SQL Server VM sizes ranging from 16 to 64 vCPU; and the HPE Hyper Converged 380 performed linearly as databases were added.

HPE Hyper Converged 380 was also able to easily consolidate a mix of Microsoft SQL Server VMs consisting of small, medium and large databases while maintaining the performance and response times necessary for transactional database workloads.

Target audience: Information in this paper will be useful for data center managers, enterprise architects, deployment / implementation engineers, Microsoft SQL Server engineers, DBAs, and others wishing to learn more about virtualizing Microsoft SQL Server on VMware vSphere with HPE Hyper Converged 380. A working knowledge of Microsoft SQL Server, server architecture, VMware vSphere, networking, and storage is recommended.

Document purpose: This document describes the performance and scalability of Microsoft SQL Server on virtual servers hosted on an HPE Hyper Converged 380 appliance and highlights recognizable benefits to technical audiences.

Introduction

IT departments face many challenges with supporting growing workload and business demands. As business demands change, IT departments struggle to rapidly create and deploy different environments to support new workloads while balancing the complexity to maintain and monitor isolated pools of infrastructure resource sprawl. HPE Hyper Converged solutions address these issues with a simple, intuitive virtualized solution that can be managed with ease.

This white paper demonstrates the use of HPE Hyper Converged 380 to deploy Microsoft SQL Server 2016 on virtual machines. The focus of this paper is to show how the HPE Hyper Converged 380 performance scales when scaling up the SQL Server database VMs as well as when scaling out the number of SQL workloads to meet the demands of growing database requirements and increasing workloads.

Deploying Microsoft SQL Server databases that support OLTP workloads on HPE Hyper Converged 380 has the following benefits.

- Rapidly develop and deliver new services for business by easily and quickly deploying virtual infrastructure. HPE Hyper Converged solutions can be easily deployed managed and maintained. This helps in reducing the time and knowledge required to deploy and maintain the system.
- Reduce infrastructure and management complexity by consolidating multiple databases on a single platform. With the capabilities
 of faster virtualization, by providing compute and software-defined storage, it is possible to host multiple databases and thereby
 utilize the resources efficiently.
- Scale business services and virtual infrastructure on-demand by adding more HPE Hyper Converged 380 platforms as workload demands grow.



HPE Hyper Converged 380 family

The HPE Hyper Converged 380 is a configurable, scalable, agile, and highly available hyper converged virtualization system. The HPE Hyper Converged 380 delivers a simple solution stack with extended flexibility and manageability. It builds on the powerful, industry standard HPE ProLiant DL380 Gen9 server platform and is combined with VMware vSphere. Using the HPE HC 380 Management UI to add full lifecycle management, VM provisioning and updates in a single pane of glass provides a unified and global experience. The HPE Hyper Converged 380 delivers a turn-key virtualization solution for medium-size business enterprises and laaS providers.

Designed from the ground up for the software-defined data center, the HPE Hyper Converged 380 enables a standardized approach to virtual server deployment. This standardized approach is regardless of whether the system is used as a primary virtualization platform in medium-size businesses or as a dedicated resource pool for specific applications in the enterprise. Unlike many hyper converged systems in the market, the HPE Hyper Converged 380 can be customized at the time of order and will be ready for virtualized workloads in a few mouse clicks.

All hardware and software components are pre-installed and pre-integrated by Hewlett Packard Enterprise. A quick customization using the HPE HC 380 Management UI software enables faster time to value unique to the HPE Hyper Converged 380. After the initial installation, IT administrators can manage their physical and virtualized environment with HC 380 Management UI and VMware vCenter Server.

Solution overview

For this solution, the HPE Hyper Converged 380 (HC 380) used for the validation testing is a four (4) node HPE Hyper Converged appliance that offers highly available server and storage. The HPE Hyper Converged 380 comes complete with all server, storage, networking, and management tools needed to begin a deployment and can be configured to support two (2) to sixteen (16) nodes per management group. From a software perspective, the HPE Hyper Converged 380 is pre-configured with VMware vSphere 6.0 and includes API integration via the HPE OneView InstantOn for VMware vCenter plug-in to facilitate simplistic platform management and solution deployment. Figure 1 shows the HPE Hyper Converged 380 system (one node).



Figure 1. HPE Hyper Converged 380 system (one node)

With valid VMware vSphere licensing applied, the HPE Hyper Converged 380 platform serves as an ideal solution for customers looking for the rapid deployment and expansion of a wide range of end-user computing.

HPE StoreVirtual VSA plays an important role in managing storage across all nodes of the HPE Hyper Converged 380 system. An instance of HPE StoreVirtual VSA runs as a virtual machine on each node of the HPE Hyper Converged 380 system, making all the local disks into a shared pool of storage. HPE StoreVirtual VSA software uses scale-out, distributed clustering to provide a pool of storage with enterprise storage features and simple management at a reduced cost. Multiple HPE StoreVirtual VSAs running on multiple servers create a clustered pool of storage with the ability to make data highly available by protecting volumes with Network RAID. Adding more HPE StoreVirtual VSAs to the cluster grows the storage pool. With Network RAID, blocks of data are striped and mirrored across multiple HPE StoreVirtual VSAs, allowing volumes and applications to stay online in the event of disk, storage subsystem, or server failure. iSCSI connectivity on HPE StoreVirtual VSA supports the use of the storage pools by VMware vSphere 6.0 as well as other applications.



HPE StoreVirtual Adaptive Optimization is an innovative technology that greatly increases the efficient use of faster storage devices by intelligently migrating data between storage devices with different performance characteristics within a single storage system. These differing storage devices, known herein as tiers, have different speeds and costs. In HPE StoreVirtual VSA, tier 0 designates the fastest storage media (SSDs) and tier 1 designates the next tier down in speed. Adaptive Optimization detects the most frequently accessed data, and in real time migrates that data to the faster tier 0 storage, and moves infrequently accessed data to the slower tier 1 storage, which is potentially lower cost storage. The net result of this selective provisioning is to offer performance approaching pure solid state at a much lower cost. For more information, see "Adaptive Optimization for HPE StoreVirtual" under the Resources and additional links section.

The solution described in this paper used a factory installed image for VMware ESXi as the base operating system. After the operating system is installed and the system is turned on, an administrator can log in to the management VM and launch HPE OneView InstantOn to complete the deployment. Then an installation wizard for HC 380 Management UI is launched. After the installation is complete, an HPE management UI VM is created. The following figure (Figure 2) shows a high-level block diagram of the tested solution configuration.

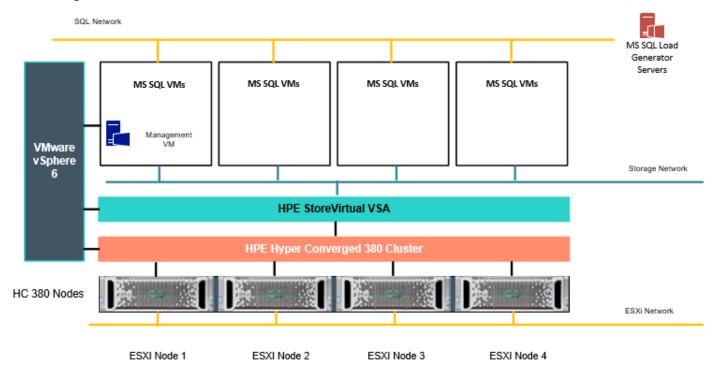


Figure 2. Logical solution diagram of HPE Hyper Converged 380 (HC 380) with Microsoft SQL Server VMs



The system is preinstalled with HPE StoreVirtual VSA VMs, a management VM, and an HPE Hyper Converged management UI VM for monitoring the entire HPE Hyper Converged 380 appliance, VM provisioning, and 3 click updates. Using VMware vSphere, SQL Server VMs are created and placed on desired server nodes. Three separate networks for management, storage, and Microsoft SQL Server VMs are defined to isolate traffic as shown in Figure 2. An external load generator server is used to drive the load. Three HPE ProLiant BL460c G7 servers are used as workload generators. The workload used in the testing for this white paper is a simulated transactional OLTP workload representative of a financial trading company. The workload generators are connected on the same network as the SQL Server VMs.

Figure 3 shows the HPE Hyper Converged 380 (HC 380) cluster with 4 nodes in VMware vSphere.

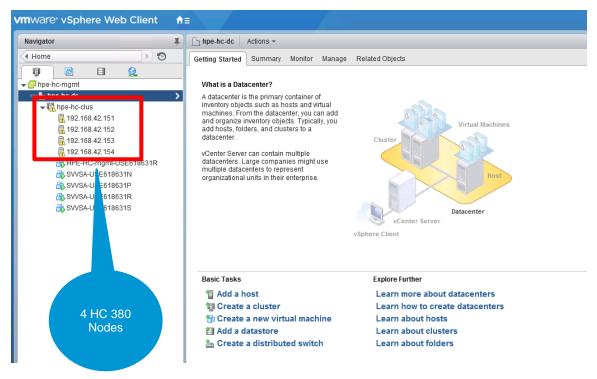


Figure 3. HPE Hyper Converged 380 cluster in VMware vSphere

Solution components

The sections below describe the hardware and software components that are used during testing.

Hardware

The HPE Hyper Converged 380 system is based on the HPE ProLiant DL380 server. HPE Hyper Converged 380 has one node per chassis, starting from two nodes and scaling up to 16 nodes in a single cluster. This solution used a single HPE Hyper Converged 380 cluster with four server nodes. All server modules have an identical node configuration with:

- Two Intel® Xeon® E5-2699 v3 processors @ 2.30GHz with 18 cores each
- 1.5 TB memory
- Two 10GbE SFP+ and four 1GbE network connectivity ports
- Hybrid storage configuration with 18 x 1.2 TB HDD and 6 x 800 GB SSDs





Figure 4. Front view of HPE Hyper Converged 380 without the bezel (one node)

Each server module has two processors of 18 cores, thus giving a total of 36 cores per server. As shown in Figure 4, each node has up to 3 storage blocks which houses a total of 24 disk drives. HPE HC 380 shared storage is provided by HPE StoreVirtual VSA and offers three types of storage blocks: Hard Disk Drive (HDD), All-Flash (SSD) and Hybrid (HDD/SSD). In this solution, each node was configured with a hybrid storage block type utilizing 3 blocks of storage containing 6 HDDs and 2 SSDs in each block. The storage on each node is configured in a local RAID set. HPE StoreVirtual VSA then creates a network RAID over these disks spanning all the nodes in the cluster (four nodes in this testing) to ensure fault tolerance in the event of any node failure. This hybrid storage configuration gives approximately 75 TB of highly available storage for a four-node system. For network connectivity, there are two 10 GbE ports on each server module, which are teamed by the operating system. Each server also has a dedicated HPE Integrated Lights-Out (iLO) management port.

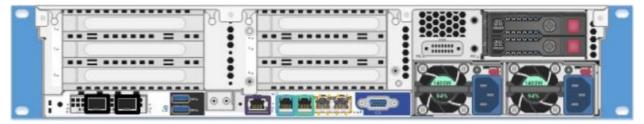


Figure 5. Rear view of HPE Hyper Converged 380

The solution made use of the HPE FlexFabric 5900AF-48XG switch (Figure 6) for networking between the management, storage, and SQL Server VM networks. The HPE FlexFabric 5900 switch series are high-density 10GbE, low latency, top-of-rack (ToR) switches. These switches are part of the HPE FlexNetwork architecture's HPE FlexFabric solution module and are ideally suited for deployments at the server access layer of large-enterprise data centers. With the increase in virtualized applications and server-to-server traffic, businesses now require ToR switch innovations to meet the needs for higher-performance server connectivity, convergence of Ethernet and storage traffic, and the capability to handle virtual environments.



Figure 6. HPE FlexFabric 5900AF-48XG switch



Software

All four nodes in the HPE Hyper Converged 380 are managed through VMware vSphere and have integrated data services from HPE StoreVirtual VSA. The system provides a virtualized, multi-tenant infrastructure based on vSphere with a single pane-of-glass management using HC 380 Management UI. A valid VMware vSphere Enterprise or Enterprise Plus license is required and may be optionally purchased from Hewlett Packard Enterprise, channel partners or an existing customer license may be used.

HPE StoreVirtual technology on top of the VMware vSphere virtualization platform provides storage capability with superior high-availability and disaster recovery capabilities. Each node on the HPE Hyper Converged 380 system has an HPE StoreVirtual VSA virtual machine running and provides a 4-node cluster with:

- Adaptive Optimization for workload acceleration (only with Hybrid Storage option)
- Network RAID 5, 6, 10, 10+1, and 10+2
- Integrated thin provisioning
- Virtual machine and application-consistent snapshots

The pre-integrated virtualization platform powered with VMware vSphere 6 enables administrators to manage their virtualized environments.

Application software

The SQL Server version used for this testing is Microsoft SQL Server 2016 (RTM) - 13.0.1601.5 (X64). This version is installed on virtual machines running Windows Server® 2016 Technical Preview 5.

Best practices and configuration guidance for the solution

The HPE Hyper Converged 380 system requires a minimum of two different networks for deployment: management and storage. For this testing, a third network is also used to connect the SQL Server virtual machines. To be able to successfully deploy the system, the following points are mandatory.

- Management and storage networks must be routable.
- VLAN must not be tagged on the management network.
- Storage network must be a tagged VLAN, to avoid collisions with management traffic.
- The third network for SQL servers must be a tagged VLAN.
- Existing DNS and AD DS domains must be running at a Windows Server 2012 R2 Datacenter functional level.

Note

For details on how to set up and configure HPE Hyper Converged 380, see "HPE Hyper Converged 380 Installation Guide" under Resources and additional links.

Three types of scenarios were tested on the HPE Hyper Converged 380 system.

- 1. Scale-up: Testing to determine how the virtualized SQL server database performance scales up when more CPU and memory are added to an individual VM
- 2. Scale-out: Testing to determine how the database performance scales out when multiple virtualized database VMs are provisioned on the HC 380 nodes.



3. Database consolidation: Testing to determine HC 380 performance characterization when scaling up the number of SQL server VMs with small and medium-size databases.

For all the test cases, the configuration for SQL Server instances were modified as documented below.

- Allow Instant File Initialization (service account given "Perform Volume Maintenance Tasks") to enable faster database file creation.
- Set Maximum Degree of Parallelism to 1 to optimize for a strictly OLTP workload.
- Increase Maximum worker threads to 3000.
- Enabled Lock Pages in Memory.
- Set Max Server Memory to an appropriate setting for the server
 - 80% of allocated virtual memory
- Enable these trace flags:
 - T834 (use large page allocations in Windows®)
 - k (set checkpoint throughput bandwidth limit)

For storage systems that are designed for low IOPS workloads, Microsoft SQL Server has an advanced startup parameter called -k, specified in Megabytes. This throttles the checkpoint I/O behavior, effectively spreading the checkpoint over a slightly longer period of time instead of a single high-bandwidth burst. By using this parameter, disk latencies maintain steady acceptable performance limits. The value of this instance parameter depends on the workload on that particular instance, and on all other concurrently running Microsoft SQL Server VMs accessing the same clustered storage system. It is important to note that spreading the checkpoint write I/O increases the database recovery time. When using this parameter it is important to ensure the recovery time meets the business SLA.

For information about the –k parameter, see the link to "SQL server checkpoint and –k startup parameter" under **Resources and additional links**.

On all servers, the following BIOS settings were enabled:

- · Hyper-Threading Enabled
- Intel Turbo Boost Enabled
- HPE Power Profile Maximum Performance

For volumes created on HPE StoreVirtual VSA:

- All volumes used were in Network RAID-10.
- Adaptive Optimization (AO) was enabled on volumes hosting SQL Server data, transaction log and tempdb files.
- Volumes created were fully provisioned.
- Multipath Extension Module (MEM) feature was installed on all the four ESXi nodes. This provides advanced multi-pathing capability for the VMware clusters. MEM reduces the disk latency with intelligent data-routing management.

For more information, see the link to "HPE StoreVirtual Storage Multipathing Deployment Guide" under Resources and additional links.

For VMware vSphere:

- Enable Jumbo frames on vSwitch for all four nodes.
- SCSI adapters were created for VM disks with paravirtual (a.k.a., PVSCSI) drivers.



· Configure VM disk with thick eager-zeroed.

The virtual machines running SQL Server were created with multiple virtual disks. In addition,

- The virtual disks of the virtual machines that were used to store SQL Server data files were of fixed type.
- The virtual disks used for operating system and SQL binaries were of dynamic type.
- Multiple volumes were provisioned on HPE StoreVirtual VSA storage. The drives for the operating system and SQL binaries were placed on the same volume. As a best practice, data and transactional files were placed on separate volumes.

Capacity and sizing

As a best practice, virtual disks must be configured across multiple virtual SCSI adapters on the virtual machine for better performance. For all test case scenarios, four SCSI adapters were used. The operating system disk drive was assigned to the first SCSI adapter. Disk drives used for data, transaction log, and tempdb were assigned to the second, third, and fourth SCSI adapters, respectively.

An OLTP style workload was used to evaluate all the test cases: scale-up, scale-out, and consolidation. Three database sizes, approximately 55 GB (small), 100 GB (medium), and 228 GB (large) were used. Databases were created with enough disk capacity to accommodate their growing size when running tests.

Scale-up scenario

The purpose of this test case was to determine how SQL Server scales up when CPU and memory resources were added incrementally. The HC 380 system with high compute capacity in terms of CPU and memory, can scale up on-demand when there is a need to satisfy higher load than usual. The space utilized by data and indexes in a VM was 800 GB and the tier 0 (SSD) capacity was 900 GB providing optimal performance under auto-tiering conditions. Table 1 shows the drive configuration for the VM used in the scale-up testing scenario. Figure 7, shows the graphical representation of the disk layout used in the scale-up test case.

Table 1. VM drives configuration for scale-up

Drive	Used for	VSA Volume Name	Adaptive Optimization	SIZE
С	Operating System	OS_VOL	Disabled	60 GB
F	Data files	SQLDATA_VOL	Enabled	800 GB
Т	Temp DB	SQLTEMP_VOL	Enabled	40 GB
L	Transaction log	SQLLOG_VOL	Enabled	40 GB



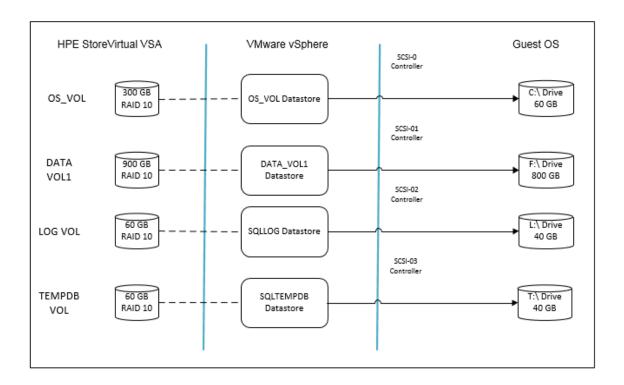


Figure 7. Disk layout for scale-up test case

During testing, CPU usage and disk latency thresholds were monitored. The load on Microsoft SQL Server was gradually increased until CPU utilization was around 90% or disk latency exceeded a 10ms threshold, thus ensuring nominal application performance to the applications connecting to the SQL Server databases. To scale the testing up, the number of simulated users was increased as the VM size increased. For the smallest virtual machine configuration, a 56 user workload was used to drive the CPU to 90%; and for the largest configuration, 264 users were simulated to drive the workload transactions. The -k parameter (described in the best practices section above) was varied to keep the disk latency under acceptable limits.

Table 2 shows the four SQL Server VM configurations, database sizes, and batch requests/sec results for the scale-up test case.



Table 2. Scale-up VM sizes

VM Core/memory	Database size	Batch requests/sec	-k value
16 vCPU/32 GB RAM	160 GB	7693	55
32 vCPU/64 GB RAM	285 GB	12447	30
40 vCPU/156 GB RAM	573 GB	16843	31
64 vCPU/256 GB RAM	688 GB	22676	32

The graph in Figure 8 shows the performance results obtained for batch requests per second on a relative scale. The graph shows that there were linear incremental gains in performance when SQL Server virtual machine resources were scaled up. The benefit of this is that administrators can anticipate system performance under increased workloads.



Figure 8. Scale-up performance

Scale-out scenario

This test case demonstrates how to scale out SQL Server instances on an HC 380 system and explores the possibility of running multiple VMs on the system without compromising the required performance needs. For this test case, a VM configuration with 32 vCPU and 64 GB memory was chosen with a database size of approximately 130 GB. Under this memory and database size configuration we have a 1:2 memory to database size ratio.

Four SQL Server VMs were created, one on each server node. The total space utilized by data and indexes was approximately 520 GB and the tier 0 (SSD) capacity is 650 GB providing optimal performance under auto-tiering conditions. An average of 56 users was used to drive the workload on each SQL Server VM. Table 3, shows the drive configuration used across all the VMs in the scale-out testing. Figure 9, shows the graphical representation of the disk layout (representing one of the four guest VMs) used in the scale-out testing.

Table 3. VM drives for scale-out test case

Drive	Used for	VSA Volume Name	Adaptive Optimization	Size
С	Operating System	OS_VOL	Disabled	60 GB
F	Data files	SQLDATA_VOL	Enabled	150 GB
Т	Temp DB	SQLTEMP_VOL	Enabled	40 GB
L	Transaction log	SQLLOG_VOL	Enabled	40 GB

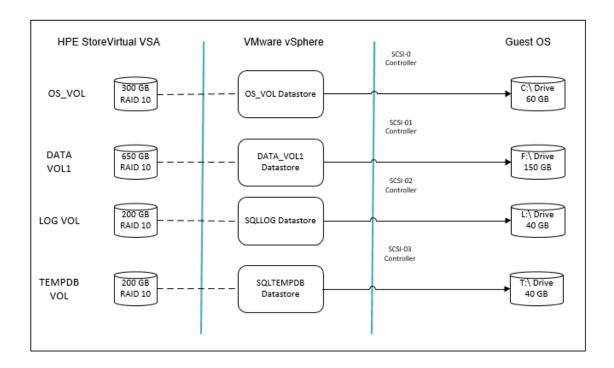


Figure 9. Single VM disk layout for scale-out test case

First, testing was started with one SQL Server VM running on one server node, and results were recorded. Next, the second SQL Server VM was started on the second node, and tests were run in parallel on the two VMs. Then, three SQL Servers and finally four SQL Server VMs were run in parallel. This way, the number of VMs was scaled out on different nodes, placing one SQL Server VM on each server node. For this scenario, with multiple SQL Server instances writing to the same disk subsystem, it was required to tune the -k parameter for the SQL Server instance running on each node based on the load. The parameter value was varied between 16 and 8.

Figure 10 shows linear performance gains as additional SQL Server VMs were introduced. With four concurrent VMs the system yielded 35,000 batch requests per second.

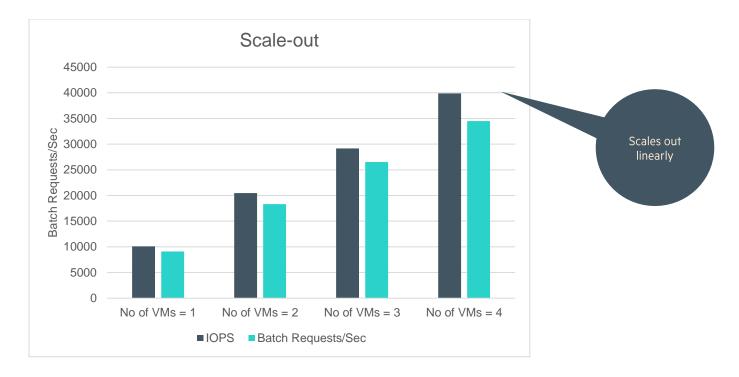


Figure 10. Scale-out performance

The HPE Hyper Converged 380 system was able to scale out easily with the disk subsystem latency well under 10ms. In this scenario, a balanced distribution of workload on the HC 380 system shows substantial gain in performance.

Consolidation scenario

In this scenario, the goal was to simulate consolidation of SQL Server instances running on multiple underutilized physical servers, onto virtualized servers, thereby reducing data center footprint and sprawl. The HC 380 has two tiers of storage, tier 0 (SSD) and tier 1 (SAS), and two different test cases were configured to highlight the performance across the tiers in a consolidation scenario. Table 4, shows the drive configuration used across all VMs and the HPE StoreVirtual VSA configuration used for the consolidation test case. Figure 11, shows a graphical representation of the disk layout used in the consolidation scenario (for a single VM).

Each VM has two disk drives for data. Tests have shown using two drives and two SCSI controllers improves performance during heavy utilization.

Table 4. VM drives for database consolidation test case

Drive	Used for	VSA Volume Name	Adaptive Optimization	Size
С	Operating System	OS_VOL	Disabled	40 GB
F	Data files	SQLDATA_VOL1	Enabled	75 GB
G	Data Files	SQLDATA_VOL2	Enabled	75 GB
Т	Temp DB	SQLTEMP_VOL	Enabled	15 GB
L	Transaction log	SQLLOG_VOL	Enabled	15 GB



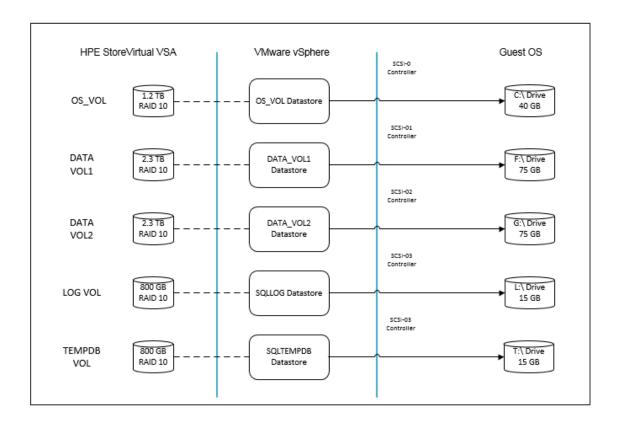


Figure 11. Single VM disk layout for consolidation test case

Tier 0 consolidation case

In the first test case, 20 VMs were concurrently deployed with all the databases provisioned in tier 0 (SSD) of the HPE StoreVirtual VSA. A mix of small and medium-size databases were used across the SQL Server instances.

Varied workloads were used across VMs, tuning them for maximum performance and an optimal value of 5MB for the -k parameter was measured. For this testing, a total workload of 88 users was driving transactions across all VMs, which resulted in a maximum batch request per second with disk latencies under the expected limit of 10ms. Table 5 lists the number of VMs and their configuration that were run on the system with acceptable latencies.

Table 5 shows an overall batch request per second transactional throughput of approximately 21,000. This is lower than the 35,000 batch request per second throughput of the scale-out test because this case used a lower (1:4) memory to database size ratio and had less SQL buffer cache resulting in more requests going to disk rather than being served from memory.

Table 5. VM configurations with utilization of tier 0

No. of VMs	VM Core/memory	Database size	Batch requests/sec	-k Value
16	8 vCPU/ 16 GB	64 GB	16572	5
4	16 vCPU/ 32 GB	114 GB	4511	5

With 20 VMs, the consolidation test in tier 0 reached 39,000 IOPS and disk latency average was less than 10ms average.



Multi-tier consolidation case

In this test case, more than 20 VMs were deployed with a combined capacity exceeding the size of tier 0, using a mix of small, medium, and large database sizes in each of the SQL Server instances. In this consolidation test a larger number of large databases were used, and the memory to database size ratio was kept the same as in the tier 0 consolidation case at 1:4.

The purpose of this test case is to show how capacity is managed as larger databases are incorporated in a tiered solution and to highlight the importance of properly sizing memory in consolidation environments.

Storage

There are two LUN options when creating an HPE StoreVirtual VSA LUN:

- LUN with Adaptive Optimization enabled
- LUN with Adaptive Optimization disabled

When a LUN has Adaptive Optimization **enabled**, writes occur on tier 0 first until tier 0 is full. After exhausting tier 0, writes occur in tier 1. As data is deemed inactive it is dynamically moved from tier 0 to tier 1.

When a LUN has Adaptive Optimization disabled writes occur directly on tier 1.

Databases with low transactional volume or low IOPS requirements should not have Adaptive Optimization enabled. Create separate LUNs and corresponding VMware datastores for small low-IOPS databases such as Development and QA databases and application metadata databases such as backup agent databases. The following diagram depicts a traditional tiered consolidation layout:

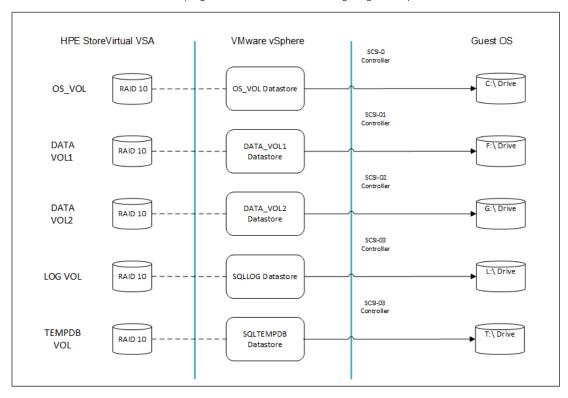


Figure 12. Single VM disk layout for database consolidation (generic)



Memory

System memory capacity and VM memory settings depend on database access patterns and combined data + index sizes. In our testing, the highest performance was achieved when memory allocation was configured with a memory to database size ratio of 1:2. Reducing the memory to a database size ratio of 1:4 resulted in more reads and slightly decreased transactional throughput.

The following diagram shows the impact of reduced memory on transactional performance and read activity. This data was measured on a 4-socket HPE ProLiant DL580 Gen9 server but is reproduced here to illustrate the impact of Microsoft SQL Server instance max memory allocation.

This data reflects a Microsoft SQL Server OLTP deployment with 8 databases having a combined data + index size of 800GB. This is not the size of datafiles, but the actual space utilized inside the datafiles.

The test starts with a generous allocation of 1.6TB of RAM (2:1 ratio), and each step lowers the SQL Server maximum memory to identify an optimum operating point in which the smallest amount of RAM yields the maximum performance.

The steps in the chart are 1.6TB / 800GB / 600GB / 400GB / 200GB / 100GB / 50GB RAM.

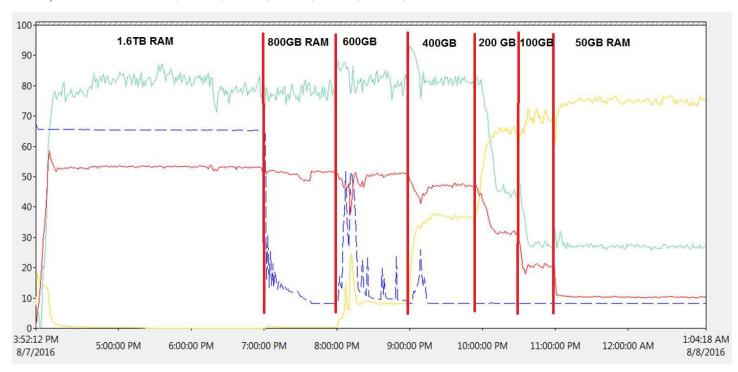


Figure 13. HPE ProLiant DL580 Gen9 memory ramp-down test

We see the optimum performance is at a 1:1 ratio (800GB), and as we lower the memory to a 1:2 ratio, reads increase (yellow line) and transactional performance is slightly reduced (red line). The impact of increased reads varies depending on the storage. Going any lower (200GB or less) decreased workload throughput significantly. We recommend performing similar tests to see if 1:2 or 1:4 works best in a particular environment.

In other words, in this particular case, a memory range of 800GB was optimal, 600GB was good and 400GB was still acceptable.

Similarly, the Multi-tier consolidation test using the HC 380 was able to accommodate multiple small medium, and several large sized databases using a 1:4 ratio. While maintaining a reasonable latency of under 10ms.



It is important to understand that the total number of SQL Server VMs supported on the HPE Hyper Converged 380 is entirely dependent on the workload, database size, and utilization. This testing is just a proof point to demonstrate scaling SQL Server VMs that are concurrently serving a reasonable workload mix.

With the Adaptive Optimization feature of HPE StoreVirtual, administrators can deploy seldom accessed databases to tier 1 by disabling Adaptive Optimization. This way, critical data such as production environment data can efficiently use tier 0 for higher performance and more efficient storage utilization in the HC 380 system.

It is critical to understand your workloads to determine the right sizing and capacity needs for your HPE Hyper Converged 380 environment. Incidentally, as you consume the resources on the HPE Hyper Converged 380, its platform extensibility allows you to easily add additional nodes to the cluster to grow your resource pools and support more Microsoft SQL Server workloads.

Analysis and recommendations

Results from the scale-up and scale-out testing demonstrate that small and medium database workloads deployed on HPE Hyper Converged 380 can scale up as well as scale out with ease. With HPE Hyper Converged 380 systems in their data centers, customers can confidently handle higher loads than usual, by increasing the virtual machine compute and storage resources. Also, the HPE Hyper Converged 380 system gives an incredible ability to virtualize and consolidate multiple traditional databases that were deployed in silos, and manage them efficiently in a single system.

When using performance as the lens for determining database or SQL instance consolidation approaches, generally speaking, databases with low utilization can be virtualized and put under one SQL Server instance with enough CPU and memory resources to satisfy the performance requirements. Multiple databases with high utilization are best put on different SQL Server instances, which ensures adequate resources for optimum performance. You should also consider on which HPE Hyper Converged 380 node the SQL Server instances and other workloads are placed on. While creating new virtual machines for new workloads, place those virtual machines on the least loaded HPE Hyper Converged 380 node.

When deploying multiple SQL Server instances, it is recommended to distribute them across all four nodes of the HPE Hyper Converged 380 system. During testing, it was observed that running multiple instances of SQL Server across different nodes demonstrated better performance and throughput when compared to running them on a single node.

The Adaptive Optimization (AO) feature in HPE StoreVirtual VSA plays a major role in providing the best performance for database workloads. Enable AO for data volumes that store SQL Server data to ensure optimum performance from the storage subsystem by storing frequently used data on the SSD tier and lesser used data on spinning media. The SSD tier is limited, and hence other workloads that do not have high throughput requirements should be stored on separate volumes with AO disabled. This way, SQL Server data can effectively make the best use of AO. During all test case scenarios, maximum performance gain was observed when the SQL Server data volume had Adaptive Optimization enabled.

Best performance was achieved when memory was allocated in a 1:2 memory-to-database ratio (or higher). When databases are deployed with a lower ratio, reads increase, and transactional throughput decreases.

Another important parameter to tune is the -k parameter in SQL Server. This parameter needs to be tuned as per the customer's environment and workload. The values used in this paper are based on the performance obtained during testing and will vary in different workload environments. In this testing, HPE Hyper Converged 380 was used only to run Microsoft SQL Server instances, but in customer environments, the HPE Hyper Converged 380 may host other application workloads as well. In such scenarios, the -k parameter must be tuned based on that particular environment.



Summary

Customer environments with IT and database sprawl can make use of the HPE Hyper Converged 380 platform to consolidate and thereby reduce both capital and operational expenditures. Multiple databases can be quickly deployed and managed on the virtual infrastructure powered by the HPE Hyper Converged 380 for Microsoft SQL Server. Testing results have also demonstrated that Microsoft SQL Server can be scaled up and scaled out with ease to handle demanding workloads.

The HPE Hyper Converged 380 system is an excellent platform for virtualizing and consolidation Microsoft SQL Server workloads without compromising on performance. It provides a combination of hardware, software, and virtualization technologies that are perfectly suited for a concurrent mix of small and medium workloads. Tiering capabilities further enable consolidation by allowing databases with different performance requirements to be co-hosted in the same system.

Resources and additional links

HPE Hyper Converged systems

hpe.com/info/hyperconverged

HPE Hyper Converged 380 system

https://www.hpe.com/us/en/product-catalog/convergedsystems/hyper-converged/pip.hyper-converged-300.1008591219.html

HPE Hyper Converged 380 Installation Guide

http://h20564.www2.hpe.com/hpsc/doc/public/display?docId=c05102727

Adaptive Optimization for HPE StoreVirtual

http://h20195.www2.hpe.com/V2/GetDocument.aspx?docname=4AA4-9000ENW

HPE StoreVirtual Storage Multi-pathing Deployment Guide

http://h20564.www2.hpe.com/hpsc/doc/public/display?docId=c05102727

HPE Reference Architectures

hpe.com/info/RA

HPE Servers

hpe.com/servers

HPE Storage

hpe.com/storage

HPE Networking

hpe.com/networking

HPE Technology Consulting Services

hpe.com/us/en/services/consulting.html

HPE Reference Architecture for Microsoft SQL Server 2016 on HPE Hyper Converged 380 system https://h20195.www2.hpe.com/V2/getpdf.aspx/4AA6-7774ENW.pdf?ver=1.0

Architecting Microsoft SQL Server on VMware vSphere – Best Practices

vmware.com/content/dam/digitalmarketing/vmware/en/pdf/solutions/sql-server-on-vmware-best-practices-guide.pdf

SQL server checkpoint and -k startup parameter

https://msdn.microsoft.com/en-us/library/ms189573.aspx



About the Authors

D.S Ravishankar, SQL Engineer, HPE
Maurice De Vidts, Lead SQL Engineer, HPE
Kannan Mani, Principal Product Manager – Data Platforms, HPE
Niran Chen Evan, Staff Solutions Architect, VMware
Deji Akomolafe, Staff Solutions Architect – Microsoft Applications Lead, VMware
Don Sullivan, Product Line Marketing Manager for Business Critical Applications, VMware

Acknowledgements

The authors thank Scott Salyer (VMware), Joe Sullivan, Don Wilson, Kimberly Kato, Said A. Syed, Rudy Castillo, Jodie Jordanger and Janet Runberg (HPE)

