Oracle 12c Databases on Hyper Converged Infrastructure using VMware vSphere 6

TECHNICAL WHITE PAPER

Contents

Executive summary

Many IT organizations have 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 HPE OneView converged management 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 Oracle database server for mid-size enterprises.

The Reference Configuration described in this paper provides customers with the expected scale-up and scale-out performance and database consolidation implications associated with deploying Oracle Database 12c OLTP workloads on HPE Hyper Converged 380.

Testing shows the HPE Hyper Converged 380 with VMware vSphere 6 can service a range of Oracle Database 12c VM sizes ranging from 4 to 16 vCPU, and the HPE Hyper Converged 380 performed linearly as databases were added. Also, when migrating a live Oracle database VM from one ESXi host to another using VMware vMotion, the database was available and consistent to users during the migration process.

HPE Hyper Converged 380 with VMware vSphere 6 was also able to easily consolidate a mix of Oracle Database 12c VMs consisting of Oracle 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, Oracle database engineers, DBAs, and others wishing to learn more about virtualizing Oracle databases on VMware vSphere with HPE Hyper Converged 380. A working knowledge of Oracle 12c Database, server architecture, VMware vSphere, networking, and storage is recommended.

Document purpose: This document describes the performance and scalability of the Oracle Database 12c on virtual servers hosted on an HPE Hyper Converged 380 appliance and highlights recognizable benefits to technical audiences.

Introduction

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 Reference Configuration demonstrates the use of HPE Hyper Converged 380 to deploy Oracle 12c databases on virtual machines using VMware vSphere 6. The focus of this paper is to show how the HPE Hyper Converged 380 performance scales when scaling up the Oracle Database 12c VMs as well as when scaling out the number of Oracle Database 12c workloads to meet the demands of growing database requirements and increasing workloads.

Deploying Oracle 12c databases that support OLTP workloads on HPE Hyper Converged 380 with VMware vSphere 6 has the following benefits:

- Rapidly develop and deliver new services for the business by easily and quickly deploying virtual infrastructure. HPE Hyper Converged solutions can be easily deployed, managed, and maintained. Also, these solutions help 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 faster virtualization and by providing compute and software-defined storage, it is possible to host multiple databases and thereby utilize the resources more 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 OneView User Experience (UX) to add full lifecycle management, VM provisioning and updates in a single pane of glass provide a unified and global experience. The HPE Hyper Converged 380 delivers a turn-key virtualization solution for medium-size business enterprises and IaaS 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 is ready for virtualized workloads in a few mouse clicks.

All hardware and software components are preinstalled and integrated by Hewlett Packard Enterprise. A quick customization using the HPE OneView User Experience (UX) 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 HPE OneView User Experience and VMware vCenter Server.

Solution overview

For this solution, the HPE Hyper Converged 380 (HC 380) used for validation testing is a four-node HPE Hyper Converged appliance that offers highly available servers 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 to sixteen nodes per management group. From a software perspective, the HPE Hyper Converged 380 is preconfigured 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).

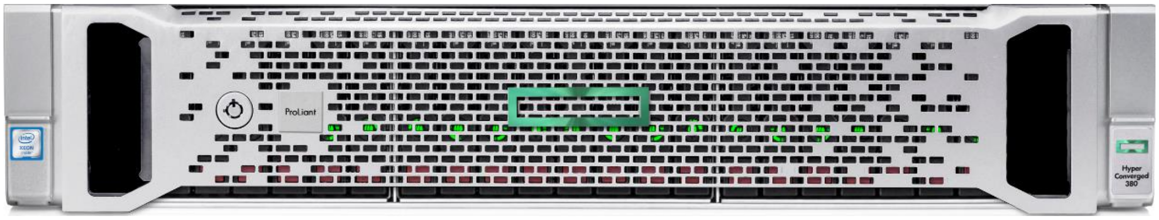

Figure 1. HPE Hyper Converged 380 system (one node)

With valid VMware vSphere licensing, 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 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 typically a lower cost storage. The net result of this selective provisioning is to offer better performance at a much lower cost.

The solution described in this paper used an HPE 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 HPE OneView User Experience is launched. Figure 2 shows a high-level block diagram of the tested solution configuration.

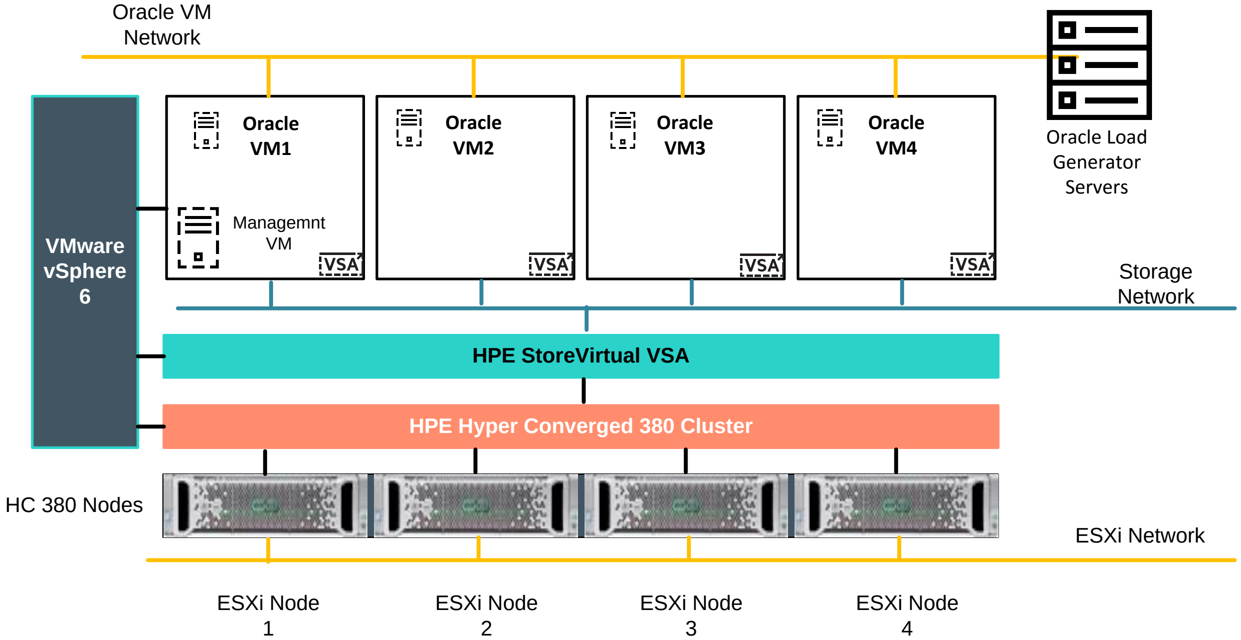

Figure 2. Logical solution diagram of HPE Hyper Converged 380 (HC 380) with Oracle 12c VMs

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, the VM provisioning, and a one-touch update for firmware and drivers. Using the VMware vSphere web client, VMs are created for running Oracle Database 12c and placed on desired server nodes. Three separate networks for management, storage, and Oracle VMs are defined to isolate traffic as shown in Figure 2. External load generator servers are used to drive the load. Six HPE ProLiant BL460c G7 servers are used as workload generators. The Oracle workload is tested using HammerDB, an open-source tool. The tool implements an OLTP-type workload (60% read and 40% write) with small I/O sizes. The workload generators are connected on the same network as the Oracle Database 12c VMs.

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

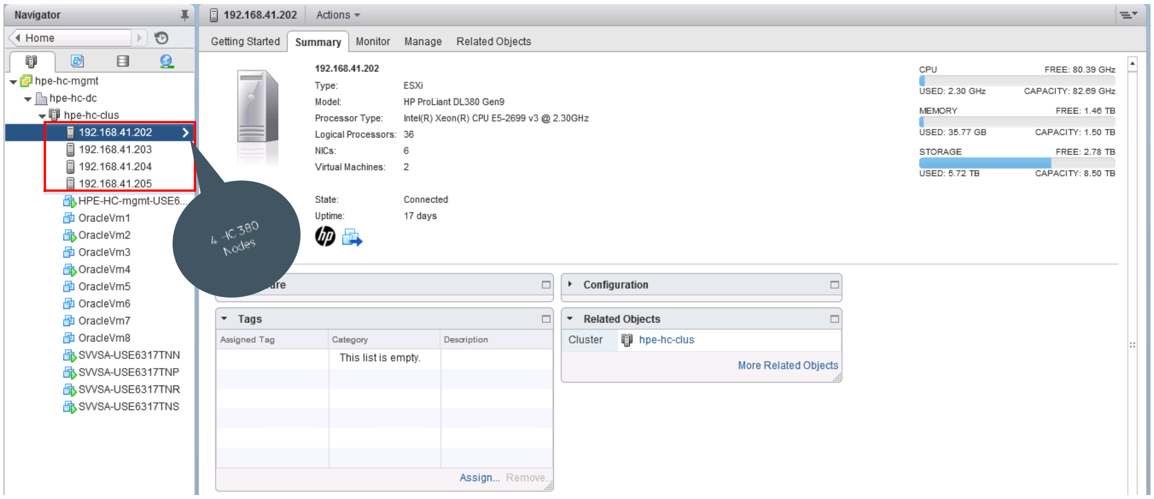

Figure 3. HPE Hyper Converged 380 cluster in VMware vSphere

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 @ 2.30 GHz with 18 cores each
- 1.5 TB memory
- Two 10 GbE SFP+ and four 1 GbE 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 18-core processors, thus giving a total of 36 cores per server. As shown in Figure 4, each node has up to three storage blocks, which house 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 three blocks of storage containing six HDDs and two 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. Using this hybrid storage configuration provides approximately 75 TB of highly available storage on our four node cluster. For network connectivity, two 10 GbE ports are available 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 used the HPE FlexFabric 5900AF-48XG switch (Figure 6) for networking between the management, storage, and Oracle VM networks. The HPE FlexFabric 5900 switch series are high-density 10 GbE, 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 of high-performance server connectivity, convergence of Ethernet and storage traffic, and 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 single pane-of-glass management using HPE OneView User Experience (UX). A valid VMware vSphere Enterprise or Enterprise Plus license is required and may be optionally purchased from Hewlett Packard Enterprise or 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, which 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
- Multi-site HA (synchronous replication across several locations)
- Remote copy (snapshot-based, asynchronous replication with bandwidth throttling)
- Storage federation with Hewlett Packard Enterprise storage products

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

Application software

Table 1 summarizes the software stack for the tests.

Table 1. Software used for workload characterization

Software	Version
Operating system	Red Hat® Enterprise Linux® (RHEL) 7.2
Oracle Database	Oracle 12c Enterprise Edition
Oracle OLTP workload	HammerDB

Test configuration and test scenarios

All tests were performed on 800 GB databases residing on RAID 10 volumes. The tests included an I/O intensive OLTP test and a CPU/IO intensive database test. The environment was tuned for maximum user transactions and maximum percentage of database usage efficiency. After the database was tuned, the transactions were recorded at different user count levels. Because many workloads have varied characteristics, the goal was to maximize transactions for each of the tests.

Table 2 gives the details about disk size used for each Oracle VM in the test. All the VMs used in the tests are configured with the same operating system, data, and redo log disk size. Also, we have ensured that the database size is the same across all VMs.

Table 2. VM disk configuration for tests

Datastore	VSA volume name	Adaptive Optimization	Size
Operating system	OS_Vol	Enabled	250 GB
Oracle Data files	Oracle_Data	Enabled	500 GB
Oracle Redolog file	Oracle_Redo	Enabled	400 GB

Figure 7 shows the disk layout for each Oracle VM used for this testing. Separate datastores have been created for the operating system, the Oracle data, and the Oracle redo log as suggested by the "Oracle Databases on VMware Best Practices Guide".

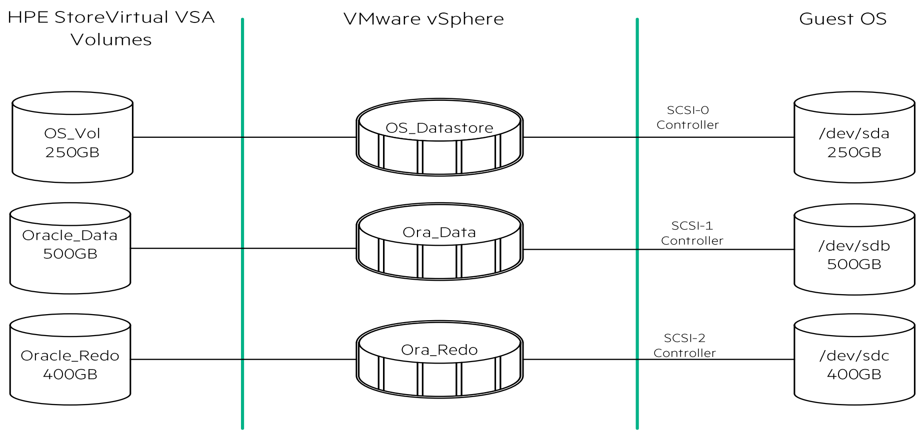

Figure 7. Disk layout for each Oracle database VM

Figure 8 shows how datastores are assigned to different SCSI controllers to get better performance. OS, Data, and Redo volumes have been assigned to SCSI controllers 0, 1, and 2 respectively.

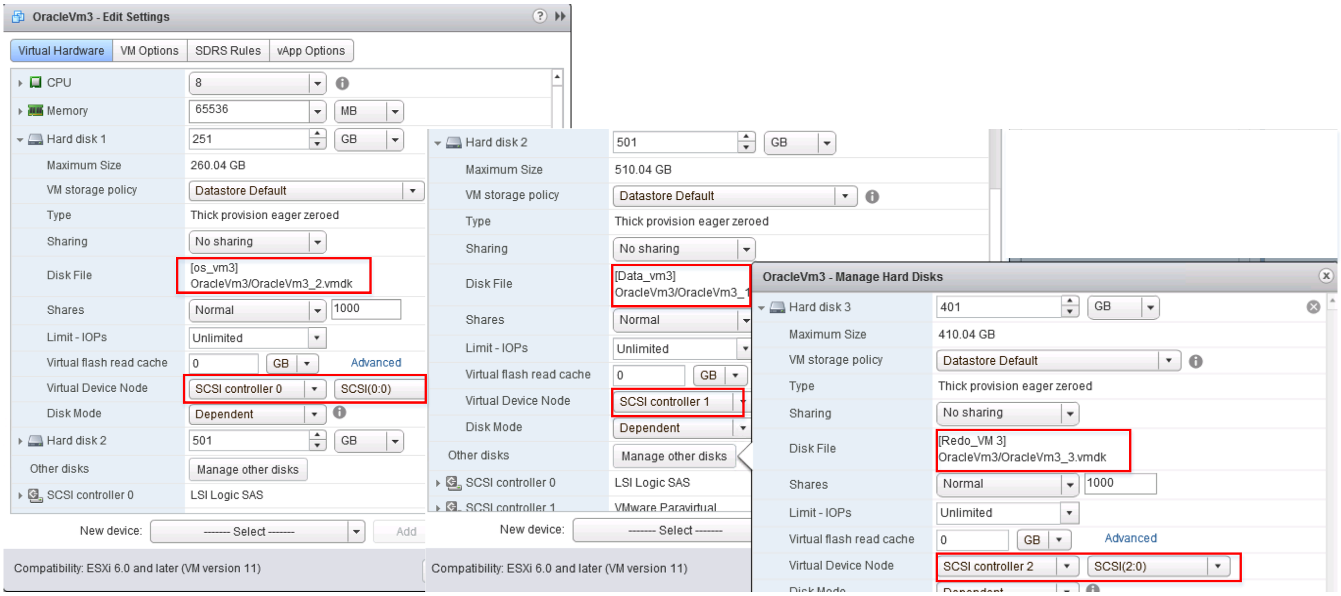

Figure 8. Datastores assigned to different SCSI controllers

For more information on VM configuration, see Appendix B.

The list below describes the tests and scenarios.

- Scale-up: Testing to determine how the virtualized Oracle Database 12c performance scales up when more CPUs and memory are added to an individual VM.
- Scale-out: Testing to determine how the database performance scales out when multiple virtualized database VMs are provisioned on the HC 380 nodes.
- Database consolidation: Testing to determine HC 380 performance characterization when scaling up the number of Oracle Database 12c VMs with medium-sized databases.
- Oracle VM migration: Impact of live migration of Oracle VM using VMware vMotion on database transactions and data consistency.

Note
All the tests were performed with Hyper-Threading off.

Table 3 summarizes the test case scenarios.

Table 3. Oracle test case scenarios

Test cases	Virtual machine (vCPU / RAM)	Oracle configuration per database (SGA in GB)	Number of Oracle users	Test case description
Scale-up single Oracle database	4 vCPU / 32 GB 8 vCPU / 64 GB 16 vCPU / 128 GB	20 GB 40 GB 80 GB	32 64 96	HammerDB workload was run against each test configuration separately, and AWR reports were captured to measure the transactions per minute.
Scale-up NUMA node aligned / non aligned	2 VMs of 16 vCPU / 128 GB 1 VM of 32 vCPU / 256 GB	80 GB 120 GB	2 x 96 144	To avoid NUMA-related performance issues, two 16 vCPU VMs were created and performance was compared with a single 32 vCPU VM on the same node.
Scale-out of Oracle databases	1 VM x 16 vCPU 2 VMs x 16 vCPU 4 VMs x 16 vCPU	120 GB 120 GB 120 GB	96 2 x 96 4 x 96	Three separate HammerDB tests were run with 1 VM, 2 VMs, and 4 VMs to measure the response time and transactions per minute. Each VM has been placed on each physical ESXi node.
Consolidation of Oracle databases	4 VMs x 8 vCPU 8 VMs x 8 vCPU	40 GB 40 GB	4 x 64 8 x 64	Two separate test cases with 4 VMs and 8 VMs were run. Number of VMs were spread evenly across the four nodes.
Oracle VM live migration using vMotion	8 vCPU / 64 GB	40 GB	64	Oracle VM was migrated from one host to another using vMotion to determine the impact of live migration on database transactions.

Capacity and sizing

For all test cases, the Oracle database size was kept the same. The Oracle workload was tested using HammerDB, an open-source tool. The tool implements an OLTP-type workload (60% read and 40% write) with small I/O sizes. The transaction results have been normalized and are used to compare the results for multiple tests. Other metrics measured during the workload come from the operating system and/or the standard Oracle Automatic Workload Repository (AWR) stats reports. For more information on the HammerDB tool, refer to hammerdb.com/hammerdb_transactionintro.pdf.

Scale-up scenario

The purpose of this test case was to determine how an Oracle Database 12c 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 to satisfy the needs of higher database workloads.

During testing, CPU usage and disk latency were monitored. The load on Oracle Database 12c was gradually increased until CPU utilization was around 90%, thus ensuring nominal application performance to the applications connecting to the Oracle 12c database. To scale up the testing, the number of simulated users was increased as the number of vCPUs were increased. For the smallest virtual machine configuration, 32 Oracle connections were used to drive the CPU to 90%. And for the largest configuration, 96 connections were simulated to drive the workload transactions.

Figure 9 depicts the performance that was achieved with each user count. The results are anchored to the 32 connection count for the 4 vCPU 32 GB VM test, which has been set to 100%. All other percentages are relative to that result. The graph shows that there were linear incremental gains in performance when resources were scaled up. Also, we have noticed that with 16 vCPU test response time was better compared to other tests. Response time increased from 4 vCPU to 8 vCPU and decreased to 6.91 milliseconds for 16 vCPU. The benefit of this response time is that administrators can anticipate system performance under increased workloads.

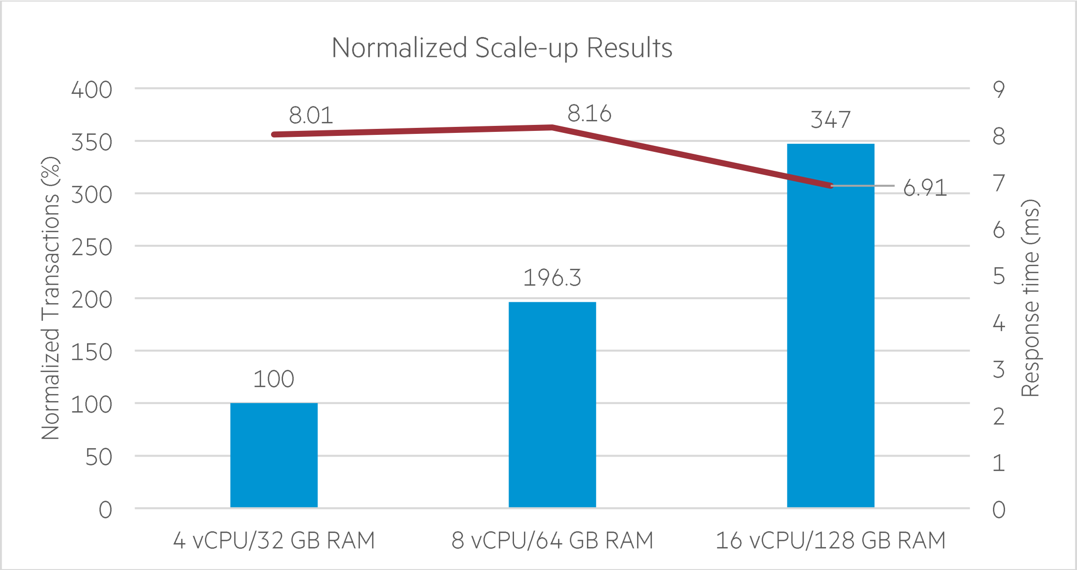

Figure 9. Scale-up performance

Each host contains two processors of 18 cores each. Because the StoreVirtual VSA VM on the nodes uses 5 vCPUs, we are left with 31 vCPUs for the guest OS. We did a scale-up test with 32 vCPU with a single VM, but the performance was not scaling linearly because 32 vCPUs span across two NUMA nodes. Figure 10 shows 32 vCPUs spanning across two NUMA nodes.

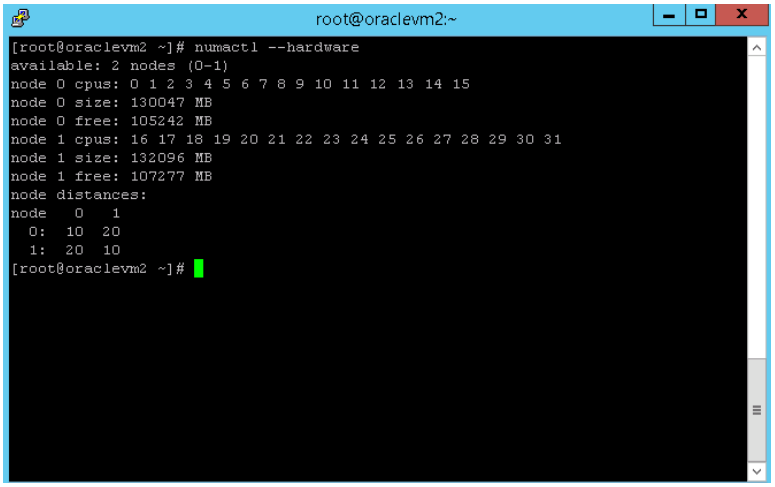

Figure 10. 32 vCPU VM spanning across two NUMA nodes

Aligning with NUMA best practices, we did a scale-up test with two 16 vCPU VMs running on the same host and compared the performance with the single 32 vCPU VM. Figure 11 shows the comparison between the scale-up test results for the single 32 vCPU versus two 16 vCPU VMs. The two 16 vCPU VMs performed 25% better than the single 32 vCPU VM in terms of transactions per minute with a lower response time, as shown in Figure 11. In addition, the dual VM configuration was able to support 33% more connections than the single VM.

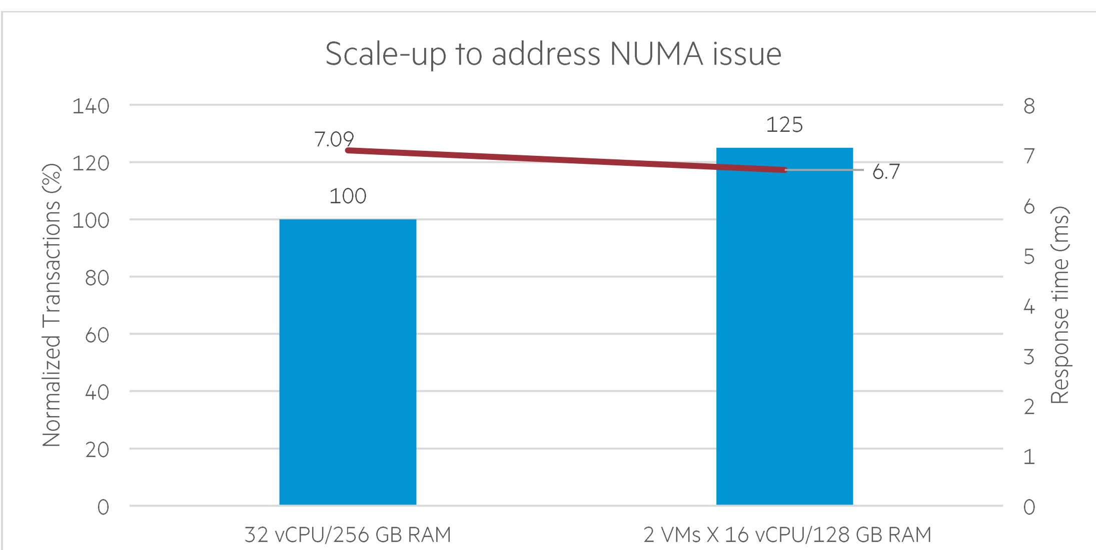

Figure 11. Scale-up results for a single 32 vCPU VM versus two 16 vCPU VMs

Scale-out scenario

This test case demonstrates how to scale out Oracle database 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, four Oracle Database 12c VMs were created, one on each host with 16 vCPUs and 128 GB memory and a database size of approximately 800 GB.

First, the testing was started with one VM running on one host, and the results were recorded. Next, the second VM was started on the second host, and tests were run in parallel on the two VMs. Then finally, four VMs were run in parallel. This way, the number of VMs was scaled out on different hosts, placing one VM on each host.

Figure 12 shows normalized transactions for the scale-out scenario with corresponding response times. Results show that the system was able to scale almost linearly from one, two, and four VMs in terms of the number of transactions per minute. Also, average response time has increased about one millisecond from one to four VM tests. This shows that the HPE Hyper Converged 380 system was able to perform better with more load with a nominal increase in response time.

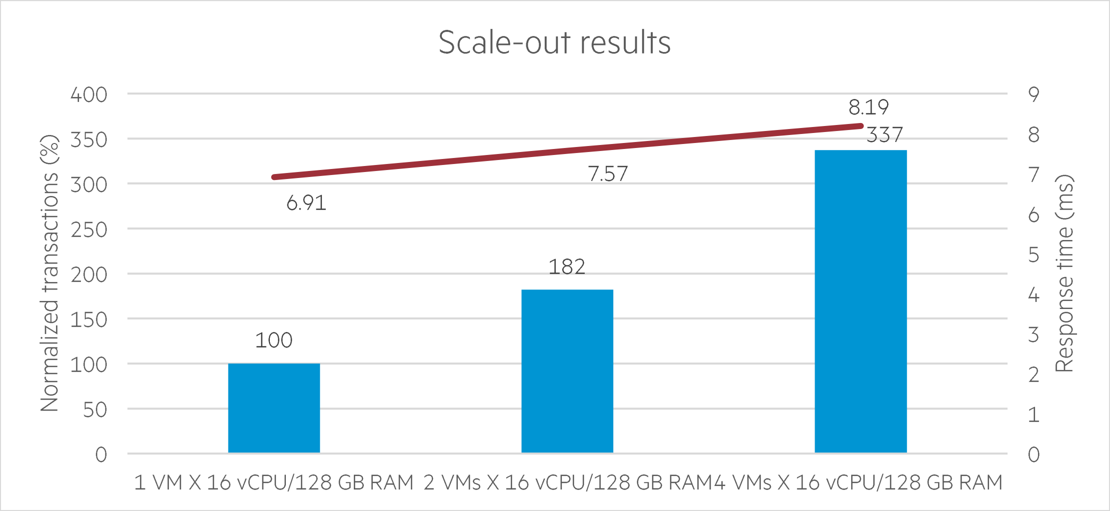

Figure 12. Scale-out performance

The HPE Hyper Converged 380 system was able to scale out easily with the disk subsystem latency well under 6 ms. 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 the consolidation of Oracle databases running on multiple underutilized physical servers onto virtualized servers, thereby reducing data center footprint and sprawl.

In this test case, 4 and 8 VMs were concurrently deployed with databases provisioned in both tier 0 (SSD) and tier 1 (SAS) of the HPE Store Virtual VSA. A HammerDB load was run in parallel on all VMs to measure the performance of the system. Figure 13 shows the consolidation results with corresponding response times.

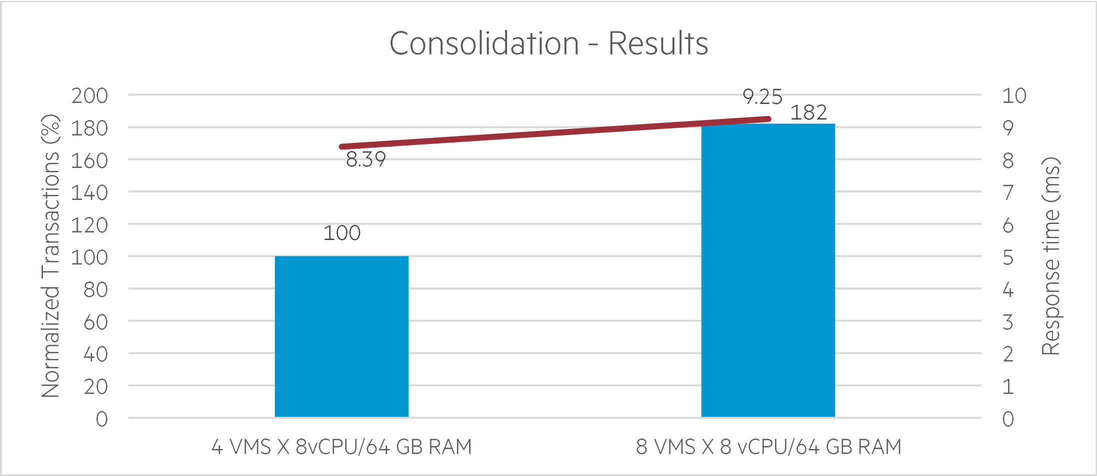

Figure 13. Normalized transactions for database consolidation case

In terms of number of transactions, the test scaled almost linearly from 4 VMs to 8 VMs as shown in Figure 13. The response time increased nominally from 8.39 to 9.25 milliseconds, which shows that the HPE Hyper Converged 380 system is able to consolidate Oracle database VMs without having much impact on transactions and response time. With the 8 VM consolidation test, we were able to reach the maximum supported StoreVirtual VSA write IOPS for an OLTP type of workload. Although HPE Hyper Converged 380 system has more compute resources, we ran out of VSA maximum supported write IOPS for OLTP workloads. The remaining compute resources on HPE Hyper Converged 380 can be used for non-OLTP types of workloads.

Live migration of Oracle database using VMware vMotion

To measure the impact of live migration of virtual machines on the Oracle database workload, we performed the following steps:

- Oracle database VM was running on host1 with the database instance started.
- HammerDB workload with 64 connections was running for approximately 20 minutes.
- Using VMware vMotion, live migration of the VM was performed moving the VM from host1 to host2.
- Impact on the OLTP workload was measured.

Figure 14 shows the impact of live migration of the Oracle VM on the OLTP workload.

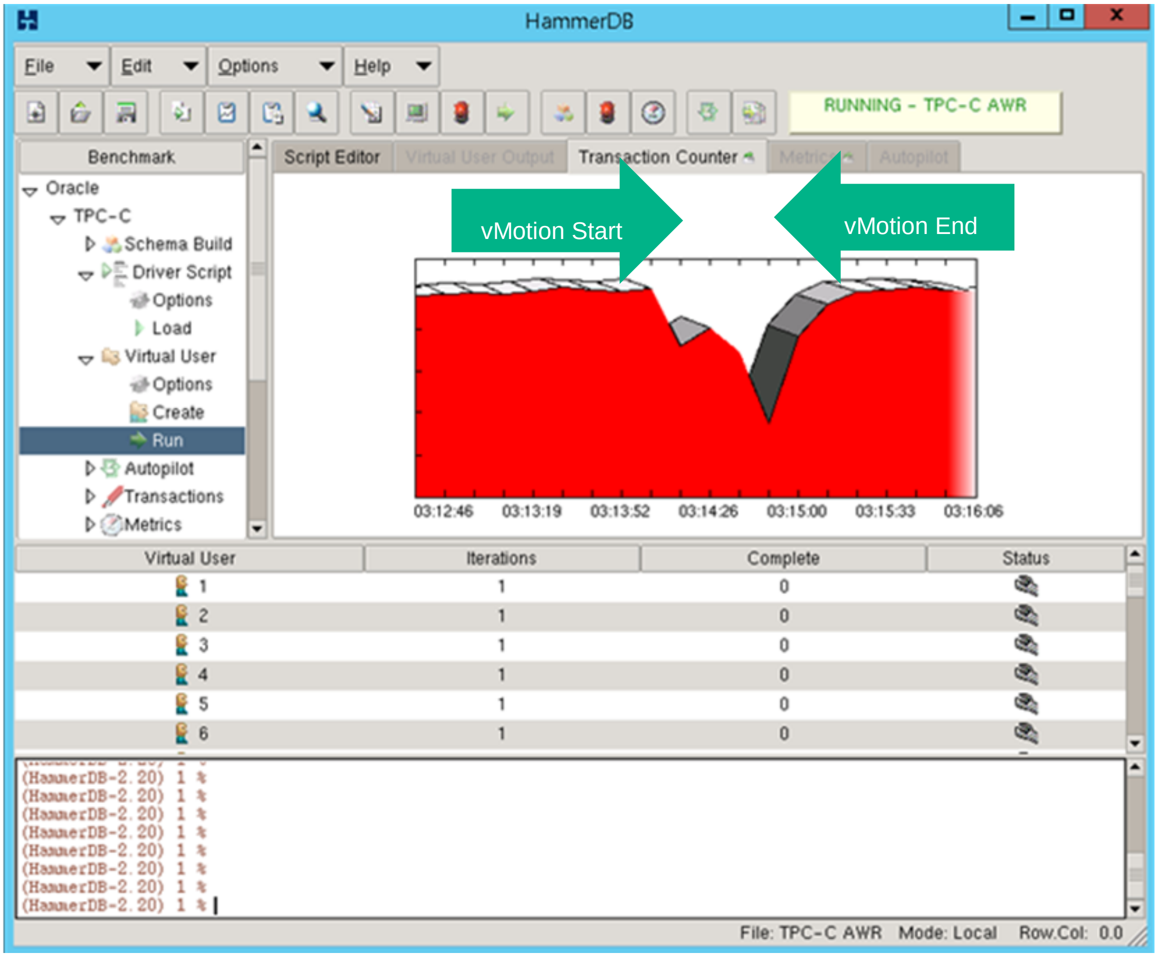

Figure 14. Impact of VMware vMotion live migration of the Oracle database VM on the OLTP workload

When the migration started, the number of transactions suddenly dipped as shown in Figure 14, and it recovered to normal level after migration was completed. It took approximately one minute to recover to a normal number of transactions. Further, there was no loss in any of the connections. The database remained intact, available, and consistent throughout the migration process. VMware administrators can use this feature to put the hosts in maintenance mode for patch updates and for other manageability actions without impacting Oracle workloads.

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 was also used to connect the Oracle Database 12c virtual machines, to separate management traffic from application traffic. To be able to successfully deploy the system, the following conditions 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 Oracle Database 12c must be a tagged VLAN to separate the management traffic from VM traffic.

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.

On all servers, the following BIOS settings were enabled:

- Hyper-Threading – Disabled
- 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 Oracle data and redo log files.
- Volumes were fully provisioned.

For VMware vSphere:

- Enable jumbo frames on vSwitch for all four nodes.
- Create SCSI adapters for VM disks with paravirtual (a.k.a., PVSCSI) drivers.
- Configure VM disk with thick eager-zeroed.
- For all disks (OS, DATA and REDO), Round Robin (RR) Multipathing policy was used and the IOPS value was set to 1.

  For more information, see https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2069356
- For more information on virtual machine configuration, see Appendix B.

For the Oracle Database:

The virtual machines running Oracle Database 12c were created with multiple virtual disks. In addition:

- Multiple volumes were provisioned on HPE StoreVirtual VSA storage. The drives for the operating system and Oracle binaries were placed on the same volume. As a best practice, data and redo log files were placed on separate volumes.
- Oracle ASM was used to manage the Oracle disk groups as a best practice.
- For all test cases, the configuration for Oracle Database 12c instances were modified as documented in Appendix C.

For RHEL 7.2:

- Tuned-adm was used to tune the RHEL kernel for best performance by setting it to throughput performance.
- Transparent huge pages were disabled.
- For all test cases, the RHEL configuration for Oracle Database 12c was modified as documented in Appendix D.

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 typical, by flexing 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 consolidation approaches, generally databases with low utilization can be virtualized with enough CPU and memory resources to satisfy the performance requirements. You should also consider on which HPE Hyper Converged 380 node the Oracle database and other workloads are placed. While creating new virtual machines for new workloads, place those virtual machines on the least utilized HPE Hyper Converged 380 node.

When deploying multiple Oracle database 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 Oracle database 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 Oracle data and redo log files to ensure optimum performance from the storage subsystem by storing frequently used data on the SSD tier and less frequently 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, Oracle data and redo log disk groups can effectively make the best use of AO. During all test case scenarios, maximum performance gain was observed when the Oracle data and redo log volumes had Adaptive Optimization enabled.

Summary

Customer environments with IT and database sprawl can make use of the HPE Hyper Converged 380 platform with VMware vSphere 6 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. Testing results have also demonstrated that Oracle Database 12c with VMware vSphere 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 consolidating database 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.

Appendix A: Oracle Virtual machine configuration

Below is the Oracle VM configuration used for this reference configuration. The same configuration was used for all Oracle VMs for all tests.

1. Create VMs with the properties as shown in Figure 15.

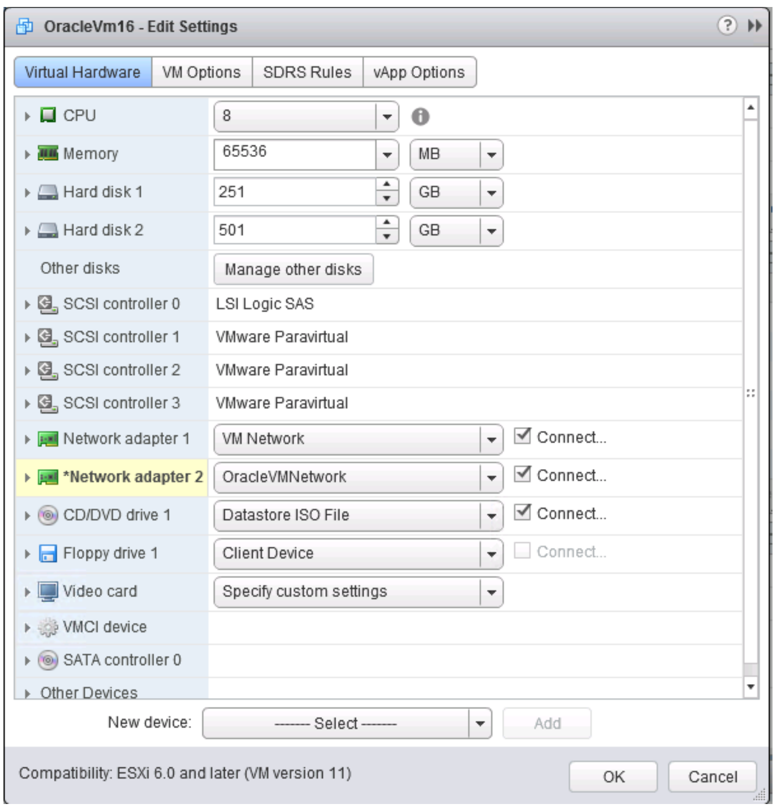

Figure 15. VMware general properties for an Oracle database VM

2. Select the **Reserve all guest memory (All locked)** checkbox in the Edit Resource Settings page as shown in Figure 16.

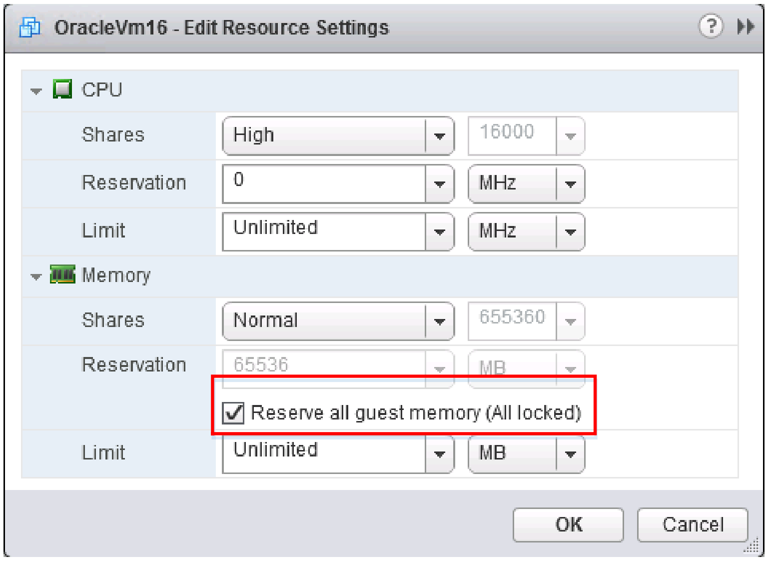

Figure 16. Reserve all guest memory to VM

3. Disable **CPU Hot Add** to make sure that NUMA nodes are managed automatically by not selecting the **Enable CPU Hot Add** checkbox as shown in Figure 17.

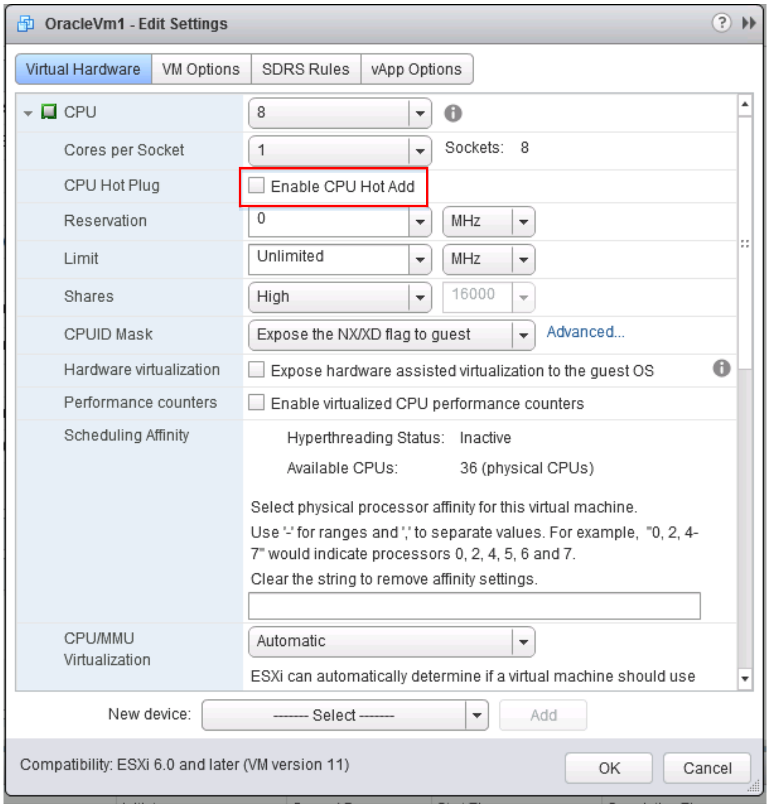

Figure 17. Disable CPU Hot Add

4. Assign the operating system, data, and redo log disks to different SCSI controllers as shown in Figure 18. Data and redo log disks are assigned to SCSI controllers 1 and 2 respectively.

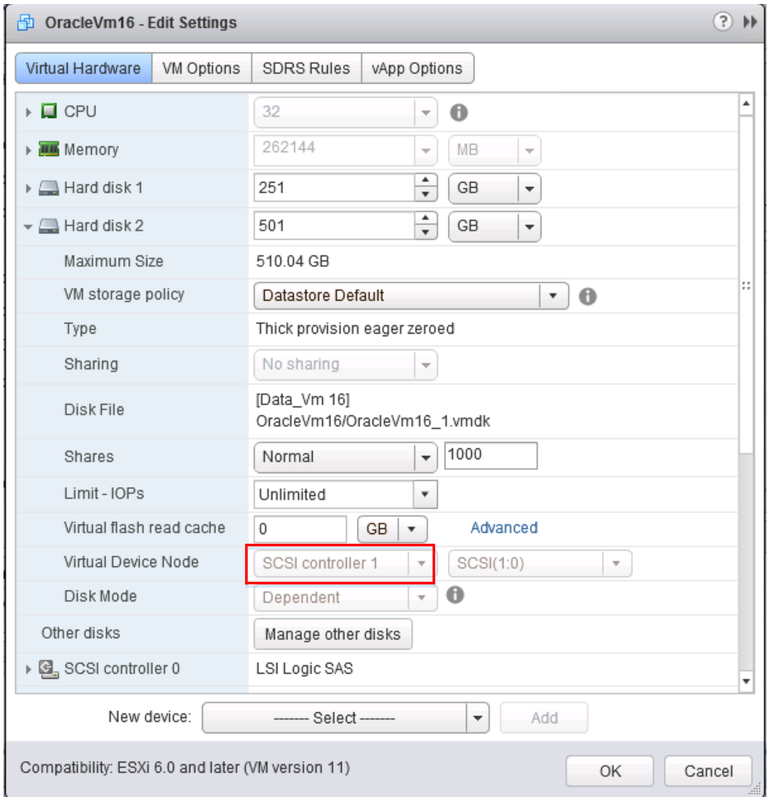

Figure 18. Assigning disks to different SCSI controllers

5. Specify the **SCSI controller 0** type as **LSI Logic SAS** for the operating system disk and **VMware Paravirtual** for data and redo disks as shown in Figure 19.

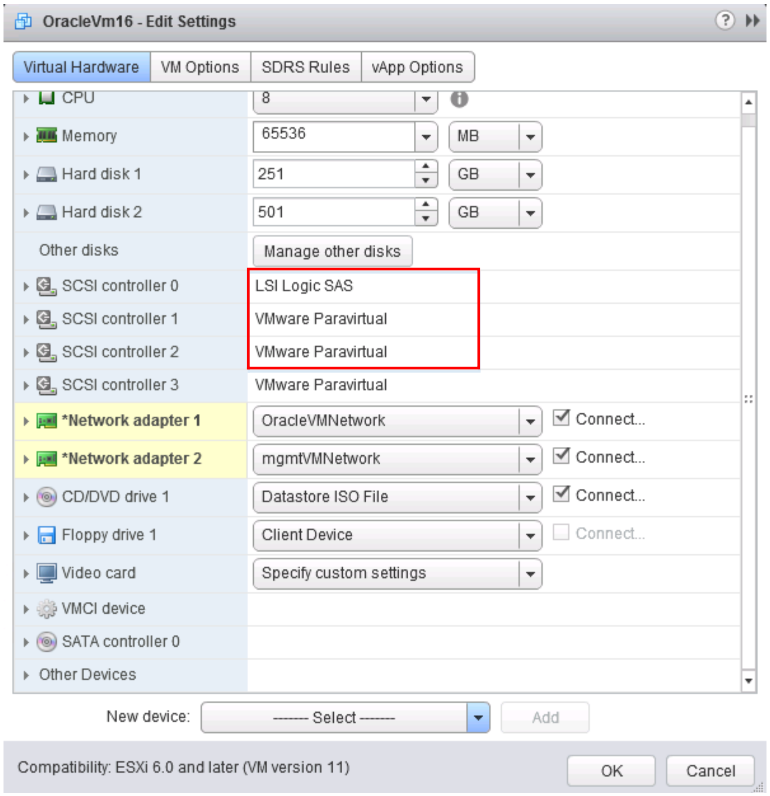

Figure 19. SCSI controller types for different disks

6. Select **VMXNET 3** from the **Adapter Type** dropdown list as shown in Figure 20.

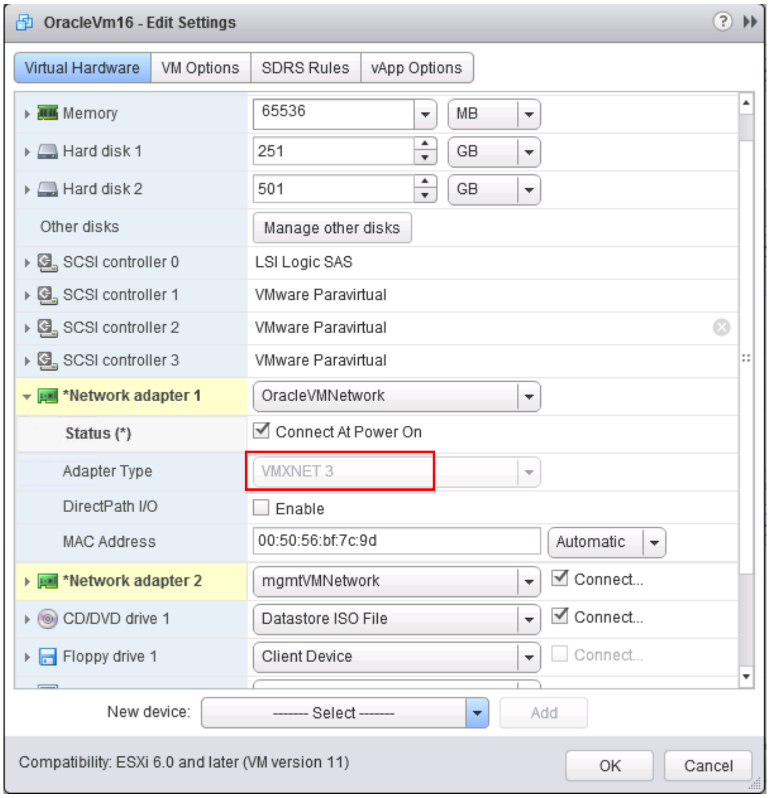

Figure 20. Selecting VMXNET 3 as network Adapter Type

7. Perform the following steps to change the disk multipathing policy to Round Robin (VMware) as shown in Figure 21.
   a. Click each ESX host, and click the **Manage → Storage** tabs.
   b. Select each disk which belongs to the VM and click on properties.
   c. Click the Edit Multipathing button.
   d. Select the policy as **Round Robin (VMware)** as shown below.

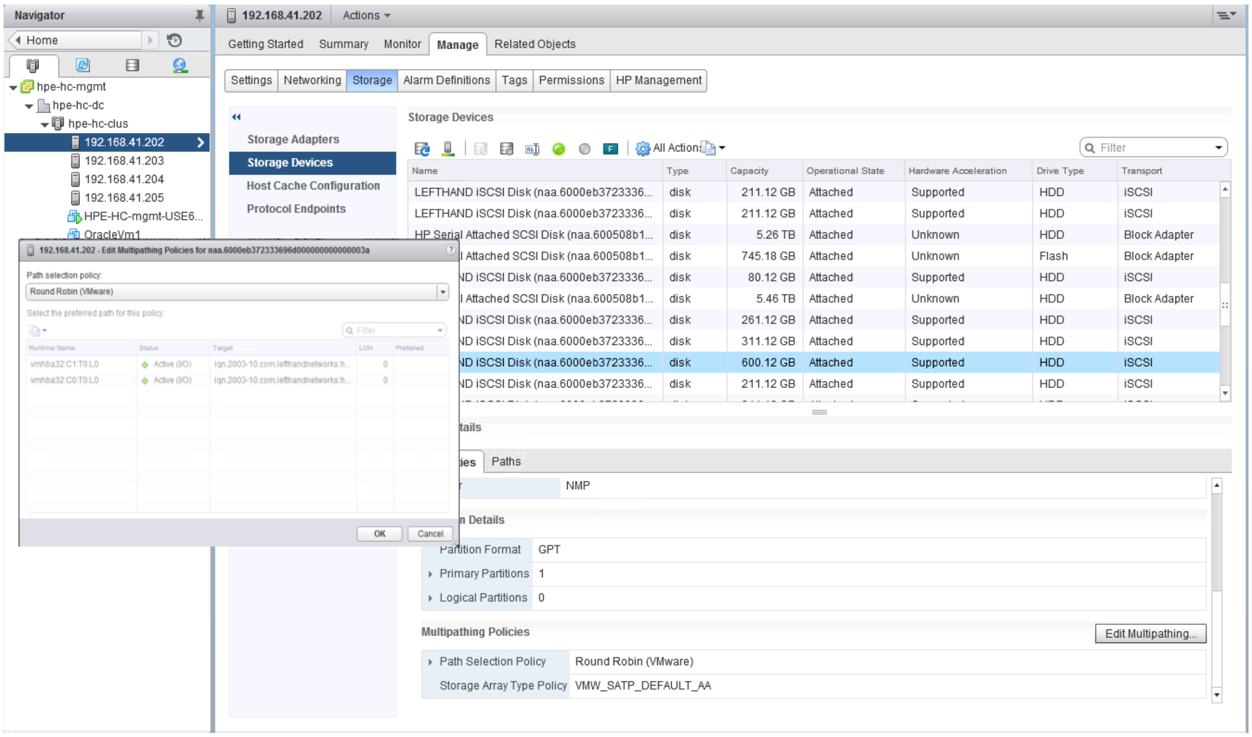

Figure 21. Changing disk multipathing policy to Round Robin (VMware)

8. Following is the script to adjust the IOPS limit to 1 for Round Robin (VMware). Execute this script on all ESX nodes after changing the disk policy to Round Robin.

```
for i in `esxcfg-scsidevs -c |awk '{print $1}' | grep naa.*`; do esxcli storage nmp psp roundrobin
deviceconfig set --type=iops --iops=1 --device=$i; done
```

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

Oracle Databases on VMware Best Practices Guide
vmware.com/content/dam/digitalmarketing/vmware/en/pdf/solutions/vmware-oracle-databases-on-vmware-best-practices-guide.pdf

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com

About the Authors

HS Siddalinga Aradhya, Oracle Engineer, HPE
Bob Augustine, Lead Oracle Engineer, HPE
Karen Dorhamer-Fadden, Lead Oracle Engineer, HPE
Kannan Mani, WW Principal Product Manager/Architect – Enterprise Data Platforms, HPE

Sudhir Balasubramanian, Staff Solutions Architect, 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)