Benchmarking Guide
June 2022

Cost-Efficient Oracle OLTP with VMware and Tiered Memory

HammerDB on Oracle Database and VMware vSphere/vSAN with Intel® Optane™ persistent memory in Memory Mode

Author

John Hubbard
Cloud Solutions Architect

Abstract

Intel® Optane™ persistent memory (PMem) is a radical technology that diversifies the memory bus, providing options to balance performance, capacity and cost. In Memory Mode, Intel Optane PMem transforms the bus into a cached architecture, using smaller DRAM footprints for cache, while Intel Optane PMem meets the capacity needs. This process is known as hardware-based memory tiering in the VMware environment and is entirely transparent to the virtual machines (VMs). For more information on deploying Intel Optane PMem in VMware environments, please refer to the Best Practices guide.

Online transaction processing (OLTP) is a database workload that most commonly represents e-commerce or ordering system usage models. Users interact or transact with a front end that generates changes to data on the back-end database(s). HammerDB TPROC-C is a synthetic benchmark derived from TPC-C to simulate OLTP workloads. HammerDB has garnered a large community over time thanks to its straightforward, easy-to-use and open-source approach.

This document provides guidance for setting up and running the HammerDB benchmark on Oracle Database within a VMware vSphere and VMware vSAN environment, to replicate Intel's performance testing and results (see the solution brief, "Reduce Memory Costs and Keep High Performance of Oracle Databases Running on VMware ESXi and vSAN"). Follow the step-by-step instructions to adhere to the same methodology used in Intel's benchmark tests. Note that this guide is not definitive for benchmarking or configuring this environment and serves as a reference for best practices.

Contents

1 Server Configuration

The benchmarking environment is composed of five servers and divided into two categories, test and support. See Section 3, "Testing Topology," for a visual description of the hardware configuration. Some dependencies exist outside the standard server hardware to achieve similar benchmark results:

- A minimum of 25 GbE network between vSAN nodes is required (this testing used 100 GbE).
- Remote Direct Memory Access (RDMA) RoCEv2 support between vSAN nodes is required.
- Jumbo frames with a Maximum Transmission Unit (MTU) of 9000 between vSAN nodes are required.
- Dynamic Host Configuration Protocol (DHCP) server on the vSAN network (if using DHCP).
- A minimum of a 1 GbE network between vSAN nodes and the support environment.
- 170,000 front-end 8K write IOPS (340K IOPS on the vSAN back end with failure to tolerate set to one).

1.1 Test Environment

The test environment is built on a 4-node vSAN cluster. A single host (target node) in the cluster is used to test different memory configurations. This host provides CPU, memory and network resources to the VMs, while the vSAN cluster provides storage.

1.1.1 vSAN Cluster

Each host is provisioned with 10 GbE and 100 GbE physical network adapters. 10 GbE networking carries HammerDB traffic from the support infrastructure and vSphere traffic for VMware vCenter and VMware ESXi. 100 GbE networking with RDMA carries vSAN and vMotion traffic. Although 25 GbE networking would suffice, 100 GbE Intel® 810 network adapters support RoCEv2 RDMA and remove all uncertainty of potential bottlenecks that could impact storage throughput. Note that Distributed Resource Scheduler (DRS) is not enabled and vMotion does not occur during testing. The tables below focus on the base hardware configuration details for the 4-node cluster and the different memory configuration comparisons.

Table 1. Server Hardware Configuration Details – vSAN Cluster

Component	Hardware Configuration
Operating System	VMware ESXi, 7.0.3, Build 18644231[a]
Platform	Intel® Server Board M50CYP
# of Nodes	4
# of Sockets	2
CPU	Intel® Xeon® Platinum 8358 processor @ 2.6 GHz
Cores per Socket/Threads per Socket	32/64
Microcode	0x0d000311
Intel® Hyper-Threading Technology	ON
Intel® Turbo Boost Technology	ON
BIOS Version	SE5C620.86B.01.1.0004
Intel® Optane™ PMem Firmware Version	2.2.0.1533
PMem BIOS Setting	Balanced
PMem Configuration Details	Memory Mode
System DDR Memory Configuration: Slots/Capacity/Speed	Variable (see table 2)
System Intel Optane PMem Configuration: Slots/Capacity/Speed	Variable (see table 2)
Total Memory per node (DRAM, PMem)	128 GB minimum
Storage	2x Intel Optane SSD P5800X 400 GB 6x SSD D7-P5510 3.84 TB
Network Interface Card	1x Intel® Ethernet Adapter 810-C 100 GbE 1x Intel Ethernet Adapter x540-AT2 10 GbE

[a] Build 18644231 was used to execute the testing referenced in this document. At the time of writing, the 7.0 Update 3 Build 18644231 is no longer available for download on the VMware website. Later builds of 7.0 Update 3 should provide no differences in results but were not used to create this document.

1.1.2 Target Node Configuration

Configure the target node using DRAM only as a baseline to perform the benchmarking. When finished, reconfigure the target node to a tiered configuration and repeat the benchmark.

Table 2. Memory Configurations

Component	1 TB DRAM	1 TB Tiered	2 TB DRAM	2 TB Tiered
DRAM to Intel® Optane™ PMem ratio	n/a	1:4	n/a	1:4
Total DRAM Capacity (GB)	1024	256	2048	512
# of DRAM DIMMs	32	16	32	16
DRAM DIMM Capacity (GB)	32	16	64	32
DRAM Speed (MT/s)	3200	3200	3200	3200
Total Intel Optane PMem Capacity (GB)	0	1040	0	2080
# of Intel Optane PMem DIMMs	0	16	0	16
Intel Optane PMem DIMM Capacity (GB)	0	128	0	128
Intel Optane PMem Speed (MT/s)	0	3200	0	3200

1.1.3 DIMM Population Rules

The memory controllers in the 3rd Generation Intel® Xeon® Scalable processors require the prescriptive placement of memory modules when operating in Memory Mode. For tiered configurations, verify that DRAM DIMMs occupy the first slot and Intel® Optane™ DIMMs occupy the second slot for each channel on the server board. For additional details, please consult the Best Practices guide.

1.2 Support Environment

The support environment is a single server running ESXi. A VM provisioned with Windows runs the HammerDB application. HammerDB connects to the VMs in the test environment via a 10 GbE physical network.

This environment does not need to be as robust as the test environment since it is not under test. For example, there is less memory and fewer CPU cores. Direct attached storage with a local Virtual Machine File System (VMFS) datastore supports the VMs versus shared storage. Simply put, this host only needs to be "good enough" to execute the workload. However, during the schema build phase of HammerDB, more vCPUs for the Windows VM reduces the time to build for each database.

Table 3. Server Hardware Configuration Details – Support Infrastructure

Component	Hardware Configuration
Platform	S2600WFT
# of Nodes	1
# of Sockets	2
CPU	Intel® Xeon® Gold 6142 processor @ 2.6 GHz
Cores per Socket/Threads per Socket	16/32
Intel® Hyper-Threading Technology	ON
Intel® Turbo Boost Technology	ON
Total Memory per Node	192 GB
Storage	1x SSD S3510 800 GB
Network Interface Card	1x Intel® Ethernet Adapter X722 10 GbE

Table 4. Windows VM and Guest Configuration

Component	Description
VM Hardware Version	19 (ESXi 7.0U2 and later compatibility)
vCPUs	24
Memory	32 GB *(no reservation)*
Network	E1000E
Storage Controller Type	LSI Logic SAS
Virtual Disk 1	150 GB; SCSI 0:0; Boot + Data
Operating System	Windows Server 2019 Standard x64; Version 1809; Build 17763.737

2 Software Configuration

2.1 Network Configuration

Simplify networking with two physical adapters for management and transactional traffic, where each adapter connects to a different virtual switch. *Note: A minimum of two network connections for each traffic type are typically recommended for production environments. In addition, traffic types are typically segmented logically with VLANs or even physically with dedicated ports and switches. The network configuration in this guide was purpose-built to perform the benchmarking and is not recommended as-is for production environments.*

2.1.1 Virtual Switch Configuration

Each virtual switch has a single physical adapter or uplink (from the host). The default standard vSwitch is used for management and the VM network. A distributed switch is used for vSAN and vMotion traffic.

Table 5. Virtual Switch Configuration

Component	Switch 1	Switch 2
Description	vSwitch0	DSwitch1
Type	Standard switch	Distributed switch
Portgroups	Management network VM network	DPortGroup
VMkernel Ports	vmk0	vmk1
Physical Adapters/Uplinks	vmnic2 10 Gbps	vmnic0 100 Gbps
RDMA Adapters	n/a	vmrdma0 *(RoCEv2)*

2.1.2 VMkernel Adapter Configuration

Two VMkernel adapters are used in the environment. vmk0 carries management traffic on the standard virtual switch with the default MTU. vmk1 carries vMotion and vSAN traffic on the distributed virtual switch with an increased MTU of 9000 for jumbo frames.

Table 6. VMkernel Adapter Configuration

Component	Adapter 1	Adapter 2
Device	vmk0	vmk1
Network Label	Management network	DPortGroup
Switch	vSwitch0	DSwitch1
MTU	1500	9000
Management	Enabled	Disabled
vMotion	Disabled	Enabled
vSAN	Disabled	Enabled
IPv4	Static IP	DHCP
IPv6	Auto via router advertisement	Disabled

2.2 vSAN Configuration

The vSAN configuration has been balanced to provide sustained write throughput and redundancy. Data services such as space efficiency and encryption have been disabled. Please refer to the vSAN and Storage Policy tables below.

Table 7. vSAN Configuration

Component	Description
Nodes	4
Disk Groups/Node	2
Cache Disk	1x Intel® Optane™ SSD P5800X 400 GB
Capacity Disk	3x SSD D7-P5510 3.84 TB
Data Services	Disabled (Space Efficiency, Encryption)

Table 8. vSAN Default Storage Policy

Component	Description
Storage Type	vSAN
Site Disaster Tolerance	None – standard cluster
Failures to Tolerate	1 failure – RAID-1 (Mirroring)
Number of Disk Stripes per Object	1
IOPS Limits for Object	0
Object Space Reservation	Thin provisioning
Flash Read Cache Reservation	0%
Disable Object Checksum	No
Force Provisioning	No
Encryption Services	No preference
Space Efficiency	No preference
Storage Tier	No preference

Virtual Machine Configuration

The VMs have been optimized to run on the target host. By default, vSphere uses large memory pages at a 1 GB granularity for VM memory management. Use a maximum of four VMs to stress the environment. Each VM is provisioned with 20% of the host memory, totaling 80%, leaving headroom for the hypervisor, etc.

Table 9. Virtual Machine Configuration

Component	Description
VM Hardware Version	19 (ESXi 7.0U2 and later compatibility)
vCPUs	16 (1 socket)
Memory	400 GB (no reservation)
Network	1x VMXNET3 adapter
Storage Controller Type	PVSCI
# of Storage Controllers	4
# of Virtual Disks	9
Virtual Disk 1	500 GB; SCSI 0:1; Data
Virtual Disk 2	500 GB; SCSI 1:0; Data
Virtual Disk 3	500 GB; SCSI 2:0; Data
Virtual Disk 4	500 GB; SCSI 3:0; Data
Virtual Disk 5	125 GB; SCSI 0:2; Redo
Virtual Disk 6	125 GB; SCSI 0:2; Redo
Virtual Disk 7	125 GB; SCSI 0:2; Redo
Virtual Disk 8	125 GB; SCSI 0:2; Redo
Virtual Disk 9	300 GB; SCSI 0:0; Boot

2.3 Guest Configuration

Oracle Enterprise Linux (OEL) is the operating system (OS) of choice and is optimized for Oracle Database. Oracle best practices were followed.

Table 10. Guest Configuration

Component	
Operating System	Oracle Enterprise Linux (OEL) 8.3
Kernel	5.4.17-2011.7.4.el8uek.x86_64
Packages	kmod-redhat-oracleasm oracle-database-preinstall-19c oracleasm-support-2.1.12-1.el8.x86_64 oracleasmlib-2.0.17-1.el8.x86_64 (Oracle CSI login required)
SELinux	Disabled
firewalld	Disabled

2.3.1 Storage Configuration

Divide storage between the boot or OS volumes and Oracle Automatic Storage Management (ASM) partitions, which are used for database data files and Redo logs.

Table 11. Guest Storage Configuration

Name	Type	Size	Filesystem	Mountpoint
sda	disk	300 G		
sda1	partition	600 M	vfat	/boot/efi
sda2	partition	1 G	xfs	/boot
sda3	partition	252 G		
ol-root	lvm	70 G	xfs	/
ol-swap	lvm	32 G	SWAP	SWAP
ol-home	lvm	150 G	xfs	/home
sdb1	partition	500 G	oracleasm	
sdc1	partition	125 G	oracleasm	
sdd1	partition	500 G	oracleasm	
sde1	partition	125 G	oracleasm	
sdf1	partition	500 G	oracleasm	
sdg1	partition	125 G	oracleasm	
sdh1	partition	500 G	oracleasm	
sdi1	partition	125 G	oracleasm	

2.3.2 Oracle ASM Configuration

Two ASM disk groups are responsible for the database. Configure both disk groups with four disks, each with external redundancy (in this case, vSAN mirroring [RAID1]).

Table 12. Oracle ASM Configuration

Disk Group	Redundancy	Allocation Unit	Capacity	Partition	Disk Name
+DATA	External (RAID 0)	1 MB	2 TB	/dev/sdb1	DATA_01
				/dev/sdd1	DATA_02
				/dev/sdf1	DATA_03
				/dev/sdh1	DATA_04
+REDO	External (RAID 0)	1 MB	500 GB	/dev/sdc1	REDO_01
				/dev/sde1	REDO_02
				/dev/sdg1	REDO_03
				/dev/sdi1	REDO_04

2.4 Oracle Grid Configuration

Install Oracle Grid before Oracle Database. This is a requirement for Oracle Restart services.

Table 13. Oracle Grid Configuration

Component	Value
Grid version	19c
Configuration	Standalone Server (Oracle Restart)
Disk Discovery Path	/dev/oracleasm/disks
Operating System Groups	OSASM | dba OSDBA | asmdba OSOPER | n/a
Oracle Base	/u01/app/grid
Inventory Directory	/u01/app/oraInventory
Software Location	/u01/app/19c/grid
Inventory Group Name	Oinstall

2.5 Database Configuration

The database was deployed in two steps, starting with the software installation followed by creation.

Table 14. Database Configuration

Component	Value
Database Version	19c – Enterprise Edition
Oracle Base	/u01/app/oracle
Software Location	/u01/app/oracle/product/19.0.0/dbhome_1
Operating System Groups	OSDBA – dba OSOPER – oinstall OSBACKUPDBA – oinstall OSDGDBA – oinstall OSKMDBA – oinstall OSRACDBA – oinstall
Creation Mode	Advanced
Configuration Type	Single instance database
Template	General Purpose – Transaction Processing
Container Database	No
Database Files Storage Type	ASM
Database Files Location	+DATA/{DB_UNIQUE_NAME}
Oracle-Managed Files (OMF)	Yes
Multiple Redo Log and Control Files	3 Groups; +REDO
Fast Recovery Area (FRA)	Yes
Recovery Files Storage Type	ASM
Fast Recovery Area	+REDO
Fast Recovery Area Size	8 GB
Archive Log	No
Listener	Port: 1521 Oracle Home; /u01/app/19c/grid
Data Vault	No
Memory Configuration Type	Automatic Shared Memory Management
System Global Area (SGA)	320 GB
Program Global Area (PGA)	40 GB

2.6 Benchmark Configuration

A Windows VM hosts the HammerDB application. HammerDB has some prerequisites for running on Oracle and is explained in detail in later sections. This section highlights the configurations for reference.

Executing the HammerDB TPROC-C (TPC-C inspired) workload is divided into three steps: install, schema build and execution or driver.

2.6.1 Installation

Table 15. HammerDB Configuration Overview

Component	Value
Operating System	Windows Server 2019 Standard
Operating System Build	Version 1809 – 17763.737
Environmental Variables – System	ORACLE_HOME = C:\Program Files\instantclient-basic-windows.x64-19.11.0.0dbru\instantclient_19_11
Oracle Instant Client[a]	Basic Windows x64 19.11.0.0
Oracle Instant Client Install Path	C:\Program Files\instantclient-basic-windows.x64-19.11.0.0dbru\instantclient_19_11
HammerDB Install Path	c:\Program Files\HammerDB-4.2

[a] *Version 19.11.0.0 was used to execute the testing referenced in this document. At the time of writing, the 19.11.0.0 version of the Instant Client basic package is no longer available on the Oracle website. The next version, 19.14.0.0, is recommended for Oracle Database 19c but was not used to create this document.*

2.6.2 Schema

Table 16. HammerDB Schema Configuration

Component	Value
Oracle Server Name	<user defined>
System User	System
System User Password	<user defined>
TPROC-C User	tpcc
TPROC-C Password	tpcc
TPROC-C Default Tablespace	tpcctab
TPROC-C Temporary Tablespace	tpcctemp
TimesTen Database Compatible	No
Use Hash Clusters	No
Partition Tables	No
Number of Warehouses	2,500
Virtual Users to Build Schema	16

2.6.3 Execution

Executing the TPROC-C workload requires configuration in three areas: the driver, virtual users and Datagen.

Driver Configuration

Table 17. HammerDB Driver Configuration

Component	Workload Configuration
Oracle Server Name	<user defined>
System User	System
System User Password	<user defined>
TPROC-C User	tpcc
TPROC-C Password	tpcc
TPROC-C Default Tablespace	tpcctab
TPROC-C Temporary Tablespace	tpcctemp
TimesTen Database Compatible	No
TPROC-C Driver Script	Timed Driver Script
Total Transactions per User	10000000
Keying and Thinking Time	No
Checkpoint when Complete	No
Minutes of Ramp-up Time	2
Minutes for Test Duration	60
Use All Warehouses	Yes
Time Profile	No
Asynchronous Scaling	No
XML Connect Pool	No

Virtual Users

HammerDB executes the workload in parallel with virtual users.

Table 18. Options for Virtual Users

Component	Workload Configuration
Virtual Users	50
User Delay (ms)	500
Repeat Delay (ms)	500
Iterations	1
Show Output	Yes
Log Output to Temp	Yes
Use Unique Log Name	Yes
No Log Bugger	No
Log Timestamps	Yes

Datagen

The Datagen configuration was left at defaults. *Note: The "Directory for File Generation" stores the logs for each run during and upon completion of the test.*

Table 19. Datagen Configuration

Component	Workload Configuration
Number of Warehouses	1
Virtual Users to Generate Data	1
Directory for File Generation	C:\Users\ADMINI~1\AppData\Local\Temp\2

3 Testing Topology

The topology has been divided into multiple sections: overview, network and storage.

3.1 Topology Overview

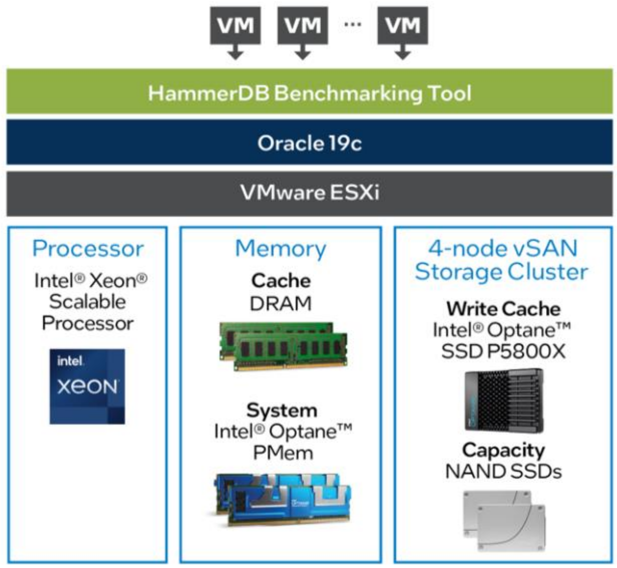

Figure 1. HammerDB connects to the database VMs, consuming memory and CPU on one host, with storage I/O to the 4-node vSAN cluster.

3.2 Host Networking Topology

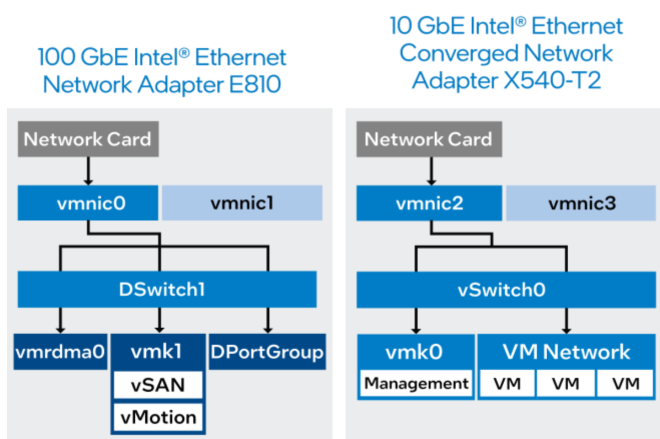

Figure 2. Management and VM network traffic utilize the 10 Gb network while vSAN and vMotion traffic utilize the 100 Gb.

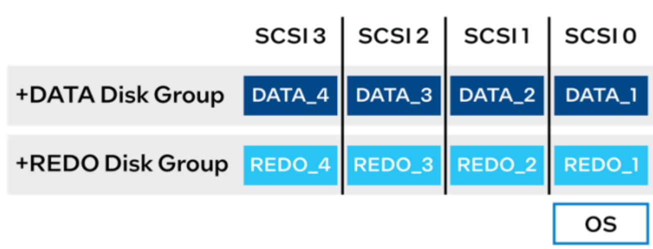

Figure 3. Virtual disks supporting the two disk groups are balanced across the four SCSI controllers to maximize throughput.

4 Recipe Details

This section provides an overview of the benchmarking process. Sections 5 through 10 provide the specific steps for each primary step, leading toward executing the actual benchmark.

4.1 Prepare the VMware Environment

The VMware environment consists of a vCenter Server Appliance (VCSA), ESXi hosts (hypervisors), vSAN and VMs. It is assumed that the reader has experience with VMware deployments; this document focuses primarily on preparing vSAN and the VMs. Deployment of the VCSA, networking and installation of ESXi are not covered in detail.

4.2 Prepare the Oracle Environment

The Oracle environment consists of the OEL OS, Oracle Grid, ASM and Oracle Database. It is assumed the reader has experience installing and configuring OEL, and while briefly covered, the primary focus is on Oracle Grid and Database configurations.

4.3 Prepare the HammerDB Environment

HammerDB can test many different databases with TPC-C or TPC-H-inspired benchmarks (TPROC-C and TPROC-H, respectively). This section covers the installing HammerDB, configuring to connect to Oracle Databases and building TPROC-C schema on Oracle Databases.

4.4 Execute the Benchmark

Execute the TPROC-C benchmark from the GUI or the CLI. The testing performed in this guide used PowerShell and the HammerDB CLI to automate repeated testing in a controlled and repeatable fashion.

5 Prepare the VMware Environment

The VMware environment consists of a VCSA, ESXi hosts (hypervisors), vSAN and VMs. It is assumed that the reader has experience with VMware deployments; the steps provided in this section focus primarily on preparing vSAN and the VMs. Deployment of the VCSA and installation of ESXi are not covered in detail.

5.1 Prerequisites

Meet the following prerequisites before configuring the vSAN cluster:

- All hosts are installed with the listed version (or newer) of ESXi.
- The VCSA is deployed with the listed version (or newer) and is reachable on the network via IP and DNS for each host.
- The VCSA manages all hosts.
- The VCSA and all hosts have Network Time Protocol (NTP) configured and enabled on startup with the same NTP server source.
- All hosts have Secure Shell (SSH) connectivity enabled.
- Physical network configured to support RoCEv2 (RDMA).
- All hosts' physical network adapters (25 Gbps or greater) are configured and enabled with RoCEv2 support.

For additional information on how to create a vSAN cluster see Appendix A: Create the vSAN Cluster.

5.2 Create the First VM

As mentioned previously, all VMs deploy on a single host. Each VM could be created independently, but it is more efficient to create one VM, then clone the configuration to additional VMs. Take note that this step only accounts for configuring the virtual hardware. The next section covers the installation of the guest OS.

The following steps summarize the process of creating the VMs. Please refer to the Virtual Machine Configuration for the complete configuration.

1. Navigate to the inventory of hosts and clusters view in the vSphere Client.
2. Right-click on the target host.
3. Select "New Virtual Machine".
4. Verify that "Create a new virtual machine" is selected and click Next.
5. Provide a name and inventory location for the VM and click Next.
6. Verify that the target host is listed as the "compute resource" and click Next.
7. Select storage:
    a. Select the "vsanDatastore".
    b. Ensure the VM Storage Policy is set to "vSAN Default Storage Policy".
    c. Click Next.
8. Verify that the VM compatibility is set to "ESXi 7.0 U2 and later" and click Next.
9. Select a guest OS:
    a. Verify that the Guest OS Family is set to "Linux".
    b. Verify that the Guest OS Version is set to "Oracle Linux 8 (64-bit)".
    c. Click Next.
10. Customize hardware:
    a. CPU
        i. Select the dropdown menu and select "16".
        ii. Click the arrow to expand the CPU section.
        iii. Verify that "Cores per Socket" is set to "16".
    b. Memory
        i. Select the text box and type "400".
        ii. Verify that the unit is set to "GB".
    c. New hard disk
        i. All capacities will be in "GB".
        ii. Select the text box and specify "500".
    d. Verify that "New SCSI controller" has "VMware Paravirtual" listed.
    e. New Network
        i. Verify that "VM Network" is selected.
        ii. Verify that the "Connect" checkbox is checked.
        iii. Expand the "New Network" section.
        iv. Verify that the "Connect at Power On" checkbox is checked.
        v. Verify that the Adapter Type is set to "VMXNET 3".
        vi. Verify that the "DirectPath I/O" "Enable" checkbox is checked.
    f. New CD/DVD Drive
        i. This setting may vary depending on your deployment.
        ii. If a copy of the guest OS is available on a datastore accessible to the host, select "Datastore ISO File" and navigate to the location. Click OK when finished.
        iii. If a copy of the guest OS is not available on an accessible datastore or is difficult to copy to one, select "Client Device".
    g. Click Next.
    h. Review the summary and click Finish to complete the wizard.

5.3 Add Storage to the VM

The initial configuration of the virtual hardware for the first VM is now complete, but additional storage is required. The steps below provide the necessary instructions to provision hard disks and SCSI controllers. To fully provision the VM storage, follow the steps below to match the storage configuration listed previously in Virtual Machine Configuration.

1. Right-click on the new VM and select "Edit Settings".

2. Click the “ADD NEW DEVICE” link in the upper-right section.
3. Select “Hard Disk”.
4. Specify 500 GB for the “New Hard Disk”.
5. Select “ADD NEW DEVICE” and select “SCSI Controller”.
6. Verify that the “New SCSI controller” is listed as “VMware Paravirtual”.
7. Select “ADD NEW Device” and select “Hard Disk”.
8. Expand the additional “New Hard Disk” section:
    a. Specify 500 GB.
    b. Select the dropdown menu next to “Virtual Device Node”.
    c. Select “New SCSI Controller”.
    d. Verify that the dropdown menu to the right lists “SCSI(1:0) New Hard Disk”.
9. Repeat Steps 6-8 until there are four SCSI Paravirtual controllers and four 500 GB hard disks:
    a. The channel or path designation will be different than selecting the new SCSI controllers for the additional hard disks in Step 8d.
    b. Verify that the third and fourth additional hard disks use the “`SCSI(2:0)`” and “`SCSI(3:0)`” paths, respectively.
10. Repeat Steps 7-8 to add four 125 GB hard disks:
    a. As mentioned in Step 9, the SCSI paths will differ for new hard disks.
    b. Verify that each additional hard disk is on a different SCSI controller.
11. Click “OK”.

5.4 Clone Additional VMs

After fully provisioning the first VM’s virtual hardware, clone the additional VMs.

1. Right-click on the new VM and select “Clone > Clone to Virtual Machine”.
2. Provide a name and inventory location for the VM and click Next.
3. Ensure the target host is listed as the “compute resource” and click Next.
4. Select “Storage”:
    a. Select the “vsanDatastore”.
    b. Verify that the VM Storage Policy is set to “vSAN Default Storage Policy”.
    c. Click Next.
5. Verify that all clone options are unchecked and click Next.
6. Click Finish to complete the wizard.
7. Repeated Steps 1-6 two more times for a total of four VMs.

6 Install and Configure Oracle Linux

Perform the install primarily using default values with some exceptions. Use a local datastore to store the ISO so that each VM can mount and install.

- Automatically provision storage with Logical Volume Manager (LVM) for the OS but specify a 32 GB swap.
- Adjust network settings to account for static IP, DNS hostname and proxy, and enable the adapter itself.
- Sync data and time settings externally.

6.1 Prepare Storage for ASM

After installing the OS, the remaining disks will need partitions before Oracle ASM can use them. Find the complete list of partitions in the previous section, Storage Configuration. The steps below outline the basic process for identifying disks and creating partitions.

1. Log in as the root.
2. Type "`lsblk`" to view the block devices.
3. Make a note of the 125 GB and 500 GB devices that do not have partitions. The devices should have names ranging from `/dev/sdb` to `/dev/sdi`.
4. For each device, use the following commands to create a partition and label, changing the target device with each iteration (sdb, then sdc, etc.):

```
parted -a optimal /dev/sdb mkpart primary 0% 100%
parted -a optimal /dev/sdb mklabel gpt
```

5. Type "`lsblk`" to view the block device and review the changes. Each device should now have a partition.

6.2 Install Support Packages

Additional packages are necessary to prepare the system to run Oracle Database. Not all of packages listed in the Guest Configuration are available in the same locations, and one requires an Oracle account to access. Install all packages with the root user.

- The "`kmod-redhat-oracleasm`" and "`oracle-database-preisntall-19c`" packages can be downloaded and installed with the DNF package manager:
```
dnf install kmod-redhat-oracleasm
dnf install oracle-database-preisntall-19c
```
- Download the "`oracleasm-support-2.1.12-1.el8.x86_64`" package from the public Oracle YUM repository (specifically "`addons x86_64`"). Then install it with the YUM package manager by performing a "`localinstall`":
```
yum localinstall oracleasm-support-2.1.12-1.el8.x86_64.rpm
```
- The "`oracleasmlib-2.0.17-1.el8.x86_64`" package requires an Oracle account to access. Once downloaded, install the package with the YUM package manager by performing a "`localinstall`":
```
yum localinstall oracleasmlib-2.0.17-1.el8.x86_64.rpm
```

6.3 Groups and Users

The "preinstall" package simplifies much of the user and group creation. Best practices recommend segmenting roles and responsibilities to manage different deployment aspects, including managing Oracle ASM. In this testing, the Oracle and Grid users are the primary accounts used to manage Oracle Database and Oracle Grid, respectively. Follow the steps below to verify and create the necessary users and groups.

1. Open a terminal or SSH session on the guest VM as root.
2. Verify the Oracle user is a member of the following groups with the "`id oracle`" command:

```
uid=54321(oracle) gid=54321(oinstall) groups=54321(oinstall),
54322(dba),54323(oper),54324(backupdba),54325(dgdba),54326(kmdba),54330(racdba)
```

3. Set the password for the Oracle user with the "`passwd oracle`" command.
Note: There is a dictionary check for password complexity, but it is not enforced. Simply type the password a second time.
4. Verify the Grid user is a member of the following groups with the "`id grid`" command:

```
uid=54331(grid) gid=54321(oinstall)
groups=54321(oinstall),54322(dba),54324(backupdba),54325(dgdba),54326(kmdba),54330(
racdba)
```

5. Create the `asmdba` group with the "`/usr/sbin/groupadd -g 54327 asmdba`" command.

6. Modify the Oracle group membership to include the new `asmdba` group:

```
/usr/sbin/usermod -u 54321 -g oinstall -G dba,asmdba,backupdba,dgdba,kmdba,racdba
oracle
```

7. Modify the Grid group membership to include the new `asmdba` group:

```
/usr/sbin/useradd -u 54331 -g oinstall -G dba,asmdba,backupdba,dgdba,kmdba,racdba grid
```

8. Set the password for the Grid user with the "`passwd grid`" command.

6.4 Configure Oracle ASM

After creating the users and groups, configure the Oracle ASM.

1. Log in as the root.
2. Type "`/usr/sbin/oracleasm configure -i`" to begin configuring the ASM library driver.
3. Enter the following:
   a. Default user to own the driver interface: `grid`
   b. Default group to own the driver interface: `asmdba`
   c. Start Oracle ASM library driver on boot (y/n) [n]: `y`
   d. Scan for Oracle ASM disks on boot (y/n) [y]: `y`
4. Wait for the Oracle ASM library driver configuration to complete.
5. Initialize Oracle ASM with the "`usr/sbin/oracleasm init`" command.
6. Note the creation of the `/dev/oracleasm` mount point.
7. Create ASM disks with the following commands. *Note: Devices associated with "`DATA`" should be the larger 500 GB disks, while "`REDO`" should be the smaller 125 GB disks. Modify the commands if the enumeration of the virtual disks is different than the [Oracle ASM Configuration.](#)*

```
/usr/sbin/oracleasm createdisk DATA_01 /dev/sdb1
/usr/sbin/oracleasm createdisk DATA_02 /dev/sdd1
/usr/sbin/oracleasm createdisk DATA_03 /dev/sdf1
/usr/sbin/oracleasm createdisk DATA_04 /dev/sdh1
/usr/sbin/oracleasm createdisk REDO_01 /dev/sdc1
/usr/sbin/oracleasm createdisk REDO_02 /dev/sde1
/usr/sbin/oracleasm createdisk REDO_03 /dev/sdg1
/usr/sbin/oracleasm createdisk REDO_04 /dev/sdi1
```

8. Scan ASM disks with the "`/usr/sbin/oracleasm scandisks`" command.
9. List ASM disks with the "`/usr/sbin/oracleasm listdisks`" command.
10. Verify ASM disks with the "`lsblk -o NAME,SIZE,LABEL,FSTYPE`" command.
11. Ensure all ASM "`DATA`" disks are 500 GB and "`REDO`" disks are 125 GB.

6.5 Disable SELinux

When set to enforcing, Security-Enhanced Linux (SELinux) prevents the installation of Oracle Grid. Disable or set SELinux to permissive to circumvent the issue.
Note: Setting to disabled will require a reboot of the system for changes to take effect, while permissive does not.

1. Log in as the root.
2. Verify SELinux status with the "`sestatus`" command.
3. If the current mode or mode from the config file is set to "`ENFORCING`", perform the following:
   a. `vi /etc/selinux/config`
   b. Press "i" to enter insert mode.
   c. Use the keyboard arrow keys to navigate to `SELINUX=ENFORCING`.
   d. Replace "`ENFORCING`" with "`PERMISSIVE`".

e. Press the ESC key.
f. Type "`:wq`" to write and quit vi. Note there is a colon before the w.

4. Restart SELinux.

6.6 Prepare Software Locations

The last steps involve setting up a well-understood directory structure for staging and operating the Oracle software. Follow these steps:

1. `mkdir -p /u01/app/19c/grid`
2. `mkdir -p /u01/app/grid`
3. `mkdir -p /u01/app/oracle`
4. `mkdir -p /u01/staging`
5. Copy the install media to the `/u01/staging` directory. There should be two archives, one for Grid and one for Database:
   a. One method is to download directly from the Oracle site via a browser on the VM to this location.
   b. Another method is to use Secure Copy Protocol (SCP) or SSH File Transfer Protocol (SFTP) to transfer the archives from another system to the VM.
6. `chown -R grid:oinstall /u01`
7. `chown oracle:oinstall /u01/app/oracle`
8. `chmod -R 775 /u01/`

7 Install Oracle Grid

Oracle Grid is the first piece of software necessary for Oracle deployment. The Grid user is the infrastructure or Grid owner, and the Grid user must perform each step of the Grid install. The installer must be executed from a virtual desktop session from the guest directly as the Grid user. Do not log in as a different user and then switch to using Grid.

Installing Grid includes the following steps:

1. Decompressing the archive.
2. Modifying the installer for OEL 8.
3. Launching the installer and creating the first ASM Disk Group for DATA.
4. Creating the second ASM Disk Group for REDO.

7.1 Decompress the Archive

Open a virtual console to the VM:

1. Log in as the Grid user.
2. Open a terminal and perform the following commands:

```
cd /u01/app/19c/grid
unzip -q /u01/staging/<name of the grid archive>
```

   (for example, grid19c3.zip)

3. Leave the terminal open.

7.2 Modify the Installer

Modify the installer to assume OEL 8.

1. From the terminal, perform the following:
    a. `vi /u01/app/19c/grid/cv/admin/cvu_config`
    b. Press "i" to enter insert mode.
    c. Uncomment `CV_ASSUME_DISTID=OEL5` (for example, remove the hashtag #).
    d. Change `OEL5` to `OEL8`.
    e. Press the ESC key.
    f. Type "`:wq`" to write and quit.
2. Leave the terminal open.

7.3 Launch the Grid Installer

From the terminal, perform the following:

1. `./gridSetup.sh`
2. Select "Configure Oracle Grid Infrastructure for a Standalone Server (Oracle Restart)" and click Next.
3. Create the ASM Disk Group:
    a. Disk Group Name: DATA
    b. Redundancy: External
    c. Allocation Unit Size: 1 MB
    d. Click the Change Discovery Path button.
    e. Set the Disk Discovery Path to `/dev/oracleasm/disks` and click OK.
    f. ASM disks should now appear.
    g. Select the DATA disks 01-04 and click Next.
4. Select "Use same passwords" for these accounts.
5. Enter the password twice to confirm and click Next.
6. Leave "Register with Enterprise Manager (EM) Cloud Control" unchecked and click Next.
7. Verify the following values:
    a. Oracle ASM Administrator (OSASM) Group: dba
    b. Oracle ASM DBA (OSDBA for ASM) Group: asmdba
    c. Oracle ASM Operator (OSOPER for ASM) Group (Optional): <blank>
    d. Click Next.
8. Click Yes to continue despite the possible invalid choice for the OSASM group.
9. Set the Oracle Base to `/u01/app/grid` and click Next.
10. Set the Inventory directory to `/u01/app/oraInventory` and click Next.
11. Check "Automatically run configuration scripts":
    a. Select "Use "root" user credential".
    b. Enter the root password.
    c. Click Next.
12. The prerequisite check will find a package, `cvuqdisk-1.0.10-1`, that has not been installed but is required:
    a. Click the Fix and Check again button.
    b. Click OK to run the Fixup Script.
13. Review the summary and click Install.
14. Observe the install progress.
15. Click Yes to run the Configuration Scripts as the root user.
16. The installer will complete.

17. Click Close to exit.
18. Leave the terminal open.

7.4 Create the Second ASM Disk Group for REDO

After installing the Grid, the ASM Configuration Assistant (ASMCA) is available. From the terminal, perform the following:

1. `cd /u01/app/19c/grid/bin`
2. Type "`./asmca`" to launch the assistant.
3. Select "Disk Groups".
4. Click the Create button:
    a. Specify the Disk Group Name as REDO.
    b. Set Redundancy to External.
    c. Set Allocation Unit Size (MB) to 1.
    d. If no disks are present, verify the Disk Discovery Path is set to `/dev/oracleasm/disks` and click OK.
    e. Select the REDO disks 01-04 and click OK.
5. There should now be two disk groups in the MOUNTED state.
6. Click Exit to close the assistant.

8 Install Oracle Database Software

Installing Oracle Database is very similar to the steps for installing Oracle Grid. All steps need to take place as the Oracle user. Installing the database software includes the following steps:

1. Decompressing the archive.
2. Modifying the installer for OEL 8.
3. Launching the installer.

8.1 Decompress the Archive

Open a virtual console to the VM. If the virtual console is still open from the Grid install, log out of the virtual desktop and log in as the Oracle user.

1. Log in as the Oracle user.
2. Open a terminal and perform the following commands:
```
cd /u01/app/oracle
unzip -q /u01/staging/<name of the database archive>
```
(for example, db19c.zip)
3. Leave the terminal open.

8.2 Modify the Installer

Modify the installer to assume OEL8.

1. From the terminal, perform the following:
    a. `vi /u01/app/oracle/product/19.0.0/dbhome_1/cv/admin/cvu_config`
    b. Press "`i`" to enter insert mode.
    c. Uncomment `CV_ASSUME_DISTID=OEL5` (for example, remove the hashtag #).
    d. Change `OEL5` to `OEL8`.
    e. Press the ESC key.
    f. Type "`:wq`" to write and quit.
2. Leave the terminal open.

8.3 Launch the Database Installer

From the terminal, perform the following steps:

1. `cd /u01/app/oracle/product/19.0.0/dbhome_1`
2. Launch the installer with the "`./runInstaller`" command.
3. Select "Set Up Software Only" and click Next.
4. Select "Single Instance Database" and click Next.
5. Select "Enterprise Edition" and click Next.
6. Specify the Oracle Base as `/u01/app/oracle` and click Next.
7. Specify the following Privileged OS Groups:
    a. Database Administrator (OSDBA) Group: `dba`
    b. All remaining groups: `oinstall`
    c. Click Next.
8. Check "Automatically run configuration scripts".
9. Select "Use "root" user credential" and enter the root password. Click Next.
10. Review the summary and click Install.
11. Click Yes to execute the Configuration scripts as the root user.
12. The installation will complete.
13. Click Close to exit.
14. Leave the terminal open.

9 Create and Configure Oracle Database

9.1 Create Oracle Database

After installing the database software, the database configuration assistant (DBCA) is available. From the terminal, perform the following steps:

1. Ensure the current working path is `/u01/app/oracle/product/19.0.0/dbhome_1`.
2. Launch the assistant with the "`dbca`" command.
3. Select "Create a database" and click Next.
4. Select "Advanced configuration" and click Next.
5. Deployment Type:
    a. Database type: Oracle Single instance database
    b. Template Name: General Purpose or Transaction Processing
    c. Click Next.
6. Database Identification:
    a. Global database name: `<fqdn>`
    b. SID: `<fqdn prefix>`
    c. Uncheck "Create as Container database".
    d. Click Next.
7. Storage Options:
    a. Select Use following for database storage attributes.
    b. Database files storage type: `ASM`
    c. Database files location: `+DATA/{DB_UNIQUE_NAME}`
    d. Select "Multiplex redo logs and control files".
    e. Type "`+REDO`" in Location for 1, 2 and 3.

   f. Click OK.
   g. Click Next.

8. Fast Recovery Option:
   a. Select "Specify Fast Recovery Area".
   b. Recovery files storage type: `ASM`
   c. Fast Recovery Area: `+REDO`
   d. Fast Recovery Area size: 8256 MB
   e. Click Next.

9. Network Configuration – Verify the Listener:
   a. Name: LISTENER
   b. Port: 1521
   c. Oracle home: `/u01/app/19c/grid`
   d. Status: Up
   e. Click Next.

10. Ensure that all Data Vault Options are unchecked and click Next.

11. Configuration Options:
    a. Specify "Use Automatic Shared Memory Management".
    b. SGA size: 320 GB
    c. PGA size: 40 GB
    d. Click Next.

12. Verify that "Configure Enterprise Manager (EM) database express" is checked with port 5500 and click Next.

13. Specify "Use the same administrator password for all accounts".

14. Provide the password twice to confirm and click Next.

15. Verify that "Create database" is the only checked option and click Next.

16. Review the Summary and click Finish.

17. Wait for the DBCA to complete.

18. Database creation is complete. Keep note of the Global Database Name and SID and Click Close to exit.

19. Leave the terminal open.

9.2 Increase Redo Log Size

By default, Redo logs are limited to 200 MB. As log files are written and encounter this limit, a log file switch occurs, moving to the next log file in the group. The following log file is marked active and becomes the new destination to record transactions. With the large number of transactions generated by the benchmark, these log file switches become increasingly frequent and can negatively impact database performance.

The primary method to reduce the negative impacts of log file switching is to increase the size of the log files. New log files must be created and added to the log group to increase the size. After adding the new log files to the group, remove the old, smaller log files to avoid reusing them in the rotation. Establish a connection to the database to make these changes. Several tools exist; however, Oracle SQL Developer is recommended for intensive development work, while SQL Plus is beneficial for quick changes. Follow the steps below to increase the size of the log file using SQL Plus:

1. Resume the open terminal session and perform the following commands.

2. Connect to the database with SQL Plus:

```
sqlplus / as sysdba
```

3. Disable word wrap for readability:

```
set wrap off;
```

4. View the current log files' group numbers and member locations on ASM:

```
select group#,member from v$logfile;
```

5. View the log group members and note the group number with the status "CURRENT":

```
select group#,status,bytes from v$log;
```

6. Add three new log files to the same directory as the existing log files. Note that the SID path will be unique for each database. In this example, groups 4-6 can be added without conflicting with the existing groups.

```
alter database add logfile group 4 ('+REDO/<SID>/ONLINELOG/group_4_2gb') size
2048m reuse;
alter database add logfile group 4 ('+REDO/<SID>/ONLINELOG/group_5_2gb') size
2048m reuse;
alter database add logfile group 4 ('+REDO/<SID>/ONLINELOG/group_6_2gb') size
2048m reuse;
```

7. Repeat Steps 4 and 5 to verify the additions and the log member currently in use.
8. Perform a manual log file switch one or more times until group 4's status becomes current as viewable in Step 4:

```
alter system switch logfile;
```

9. Verify that the log file groups 1-3 are marked as "INACTIVE." If one or more is marked as "ACTIVE," perform a manual system checkpoint:

```
alter system checkpoint global;
```

10. Remove the previous log file groups 1-3, so only the 2 GB log files remain:

```
alter database drop logfile group 1
alter database drop logfile group 2
alter database drop logfile group 3
```

11. Leave the SQL Plus Session open.

9.3 Enable ORACLE_SCRIPT

HammerDB connects to the database via the system user. The system user then creates a new user and tablespaces to store the TPROC-C schema. It is important to note that errors encountered with the HammerDB and Oracle database versions during this stage were mitigated by enabling ORACLE_SCRIPT and are not recommended for production deployments.

1. Resume the existing SQL Plus session.
2. Enable ORACLE_SCRIPT:

```
alter session set "_ORACLE_SCRIPT"=true;
```

3. Type "`quit`" to disconnect the SQL Plus session.
4. Leave the terminal open.

9.4 Retrieve Transparent Network Substrate (TNS) Names

HammerDB interacts with the Oracle Instant Client to connect to Oracle Database. The Instant Client uses TNS names to identify and connect to one or more databases. A later section discusses updating the `tnsnames.ora` file on the benchmark system. Copy the contents of the local `tnsnames.ora` files on the Oracle systems to identify the updates easily. Perform the following steps:

1. Resume the open terminal session.
2. View the contents of the `tnsnames.ora` file:

```
cat /u01/app/oracle/product/19.0.0/dbhome_1/network/admin/tnsnames.ora
# tnsnames.ora Network Configuration File:
/u01/app/oracle/product/19.0.0/dbhome_1/network/admin/tnsnames.ora
# Generated by Oracle configuration tools.

LISTENER_ORCL1 =
  (ADDRESS = (PROTOCOL = TCP)(HOST = oeldb01.example.com)(PORT = 1521))

ORCL1 =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP)(HOST = oeldb01.example.com)(PORT = 1521))
    (CONNECT_DATA =
```

```
      (SERVER = DEDICATED)
      (SERVICE_NAME = orcl1.example.com)
    )
  )
```

3. Make a note of the file and copy the contents for later use. *Note: Names in example above were altered to protect confidentiality.*

10 Prepare the HammerDB Environment

Preparing the HammerDB environment consists of the following steps:

1. Download and install HammerDB for Windows.
2. Download and configure the Oracle Instant Client.
3. Create the ORACLE_HOME environmental variable.
4. Duplicate and configure HammerDB instances.
5. Build the schemas.

Before continuing, ensure that the support infrastructure and Windows VM have sufficient network connectivity to download the required files directly or access them from a shared location. It is assumed the Windows VM has internet connectivity.

10.1 Download and Install HammerDB for Windows

1. Establish a remote desktop session (RDP) to the Windows VM.
2. Launch a web browser and navigate to Releases · TPC-Council/HammerDB · GitHub.

   Note: At the time of this writing, the most recent version was listed (4.4). Version 4.2 was used during the testing outlined in this guide. The most recent version should produce similar results but was not tested and equivalent performance between versions is not guaranteed.
3. Scroll down to Version 4.2 and click the download link for Windows Setup. A direct link has also been provided: https://github.com/TPC-Council/HammerDB/releases/download/v4.2/HammerDB-4.2-Win-x64-Setup.exe.
4. Navigate to the download location.
5. Launch the installer: `HammerDB-4.5.Win-x64-Setup.exe`.
6. Click Next.
7. Accept the agreement and click Next.
8. Accept the default Installation Directory: `C:\Program Files\HammerDB-4.2` and click Next.
9. Click Next to begin the install.
10. Wait for the install to complete.
11. Uncheck both boxes and click Finish.
12. Leave the RDP session open.

10.2 Download and Configure the Oracle Instant Client

1. Return to the web browser and navigate to https://www.oracle.com/database/technologies/instant-client/downloads.html.
2. Click the download link: Instant Client for Microsoft Windows (x64).
3. Scroll down to Version 19. *Note: At the time of this writing, the most recent version, 19.14, was listed. Version 19.11 was used during the testing outlined in this guide. The most current version should produce similar results, but was not tested and equivalent performance between versions is not guaranteed.*

4. Click the download link for the Basic Package: instantclient-basic-windows.x64-19.14.0.0.0dbru.zip.
5. Navigate to the download location.
6. Decompress or extract the entire archive to a new directory: `C:\Program Files\instantclient-basic-windows.x64-19.11.0.0.0dbru`.
7. Navigate to the new path: `C:\Program Files\instantclient-basic-windows.x64-19.11.0.0.0dbru\instantclient_19_11\network\admin`.
8. Open the `tnsnames.ora` files with a text editor.
9. Refer to TNS names that were coped from the Oracle VMs in step Retrieve Transparent Network Substrate (TNS) Names.
10. Copy and paste the contents of each Oracle VM's local `tnsnames.ora` file into the `tnsnames.ora` file for the Instant Client.
11. Save changes and exit the text editor.
12. Leave the RDP session open.

10.3 Create the ORACLE_HOME Variable

HammerDB relies on the ORACLE_HOME environmental variable when testing Oracle Database from a Windows system. This variable is not created during HammerDB install and must be done manually. To create the variable, perform the following:

1. Right-click the Windows Start menu.
2. Select System.
3. Click the System info link under Related settings in the top-right corner.
4. Click Advanced system settings.
5. Click the Environmental Variables button.
6. Under variables, click the New button.
7. Specify the name and value for the new variable. The value will be the path to the Oracle Instant Client:
   a. Variable name: ORACLE_HOME
   b. Variable value: `C:\Program Files\instantclient-basic-windows.x64-19.11.0.0.0dbru\instantclient_19_11`
   c. Click OK.
8. Click OK to exit the Environmental Variables window.
9. Click OK to exit the System Properties window.
10. Leave the RDP session open.

10.4 Duplicate HammerDB Instances

Multiple instances of HammerDB are required to benchmark multiple single-instance databases simultaneously. Copy or duplicate the HammerDB install path for each database to create replications. Number the duplicated HammerDB install paths to make it easier to track the database being tested by a specific instance. Configure HammerDB_1 to test db01, and so on. To duplicate the HammerDB instances, perform the following steps:

1. Navigate to the parent directory of the HammerDB 4.2 install: `C:\Program Files`.
2. Right-click the HammerDB-4.2 folder and select Copy.
3. Right-click in the empty space of the folder window and select Paste.
4. Right-click on the copy and select Rename.

5. Rename the directory `HammerDB-4.2_#` where the number reflects the database targeted for testing.
6. Repeat Steps 2-5 for each additional database.
7. Creating shortcuts to each HammerDB instance is recommended to save time.

10.5 Build the Schema

Configure each HammerDB instance to connect to a particular database VM, and then build the schema.

1. Navigate to the first HammerDB instance directory (for example, `C:\Program Files\HammerDB-4.2_1`).
2. Double-click `hammerdb.bat`.
3. The HammerDB GUI will launch.
4. Under Benchmark, select Oracle.
5. Expand TPROC-C.
6. Expand Datagen:
   a. Double-click Options.
   b. Note that the Directory for File Generation defaults to the current user's `\AppData\Local\#` directory. All HammerDB logs from this instance will be stored here. Feel free to change the location if necessary.
   c. Click Cancel.
7. Expand Schema Build:
   a. Double-click Options.
   b. Oracle Service Name: (for example, orcl1)
   Note: The Oracle Service Name links the HammerDB instance to the database under test via the Oracle Instant Client's `tnsnames.ora` file.
   c. For the remaining fields, refer to the Schema configuration.
   d. Click OK.
   e. Double-click Build.
   f. Click Yes to build the schema.
   g. The schema will take time to build.
8. Repeat Steps 1-7 for the remaining HammerDB instances. *Note: Depending on the resources available to the Windows VM, it may be possible to build multiple schemas concurrently. If so, pay attention to the number of virtual users assigned to build the schema. It is recommended not to exceed a 1:1 ratio between virtual users and physical cores in the support infrastructure.*

10.6 Troubleshoot

If the schema builds fail or encounter complications, the schema can be removed with a single SQL statement `drop user tpcc cascade` (where `tpcc` is the account owner for the tablespace responsible for the schema). Once the user has been dropped, the schema build can be re-attempted. *Note: For general troubleshooting, reducing the warehouses and virtual users decreases the size of the schema, thus the time to build it. This approach saves lots of time during troubleshooting.*

10.7 Preserve the Schema

Once the schema has been built on each database, the benchmark is primed and ready to be executed. However, running the benchmark is a destructive action. After building and testing the schema, it must be rebuilt before testing again. There are myriad ways to approach this requirement for multiple test iterations:

- Rebuild the schema every time.
- Backup and restore the database.
- Use Oracle Data Pump.
- Use VMware snapshots.
- Convert virtual disks to Independent Non-Persistent mode.
- Clone VMs.

Each approach has varying tradeoffs, ranging from time, performance or stability impacts, and overall complexity. During the time of testing, the best overall solution found was to clone the VMs.

Cloning the VMs for each test iteration is straightforward and consists of a few steps:

1. Shut down each VM gracefully and wait for them to become powered off.
2. Clone each VM to a new name (for example, oeldb01 > clone-oeldb01).

11 Execute the Benchmark

Running the benchmark is straightforward and can be executed from the GUI or CLI. However, running multiple iterations repeatedly across different hardware configurations requires additional steps. A "Start from Zero" approach was used to ensure consistency across test iterations and configurations. Starting from zero involves the following steps:

1. Power off and delete the cloned VMs.
2. Clone each source VM.
3. Put the host into maintenance mode (with no vSAN Data Migration) and reboot.
4. Power on each cloned VM.
5. Begin data collection if using automation; otherwise, manually export vSphere Performance metrics via the vSphere client.
6. Execute the benchmark and wait for completion.
7. Stop data collection.
8. Repeat Steps 1-7 for each iteration.

11.1 Run HammerDB

1. Navigate to the first HammerDB instance directory (for example, `C:\Program Files\HammerDB-4.2_1`).
2. Double-click `hammerdb.bat`.
3. The HammerDB GUI will launch.
4. Under Benchmark, select "Oracle 533 221 high 5".
5. Expand TPROC-C.
6. Expand Driver Script:
   a. Double-click Options and refer to Driver Configuration.
   b. Oracle Service Name: (for example, orcl1). *Note: The Oracle Service Name links the HammerDB instance to the database under test via the Oracle Instant Client's `tnsnames.ora` file.*
   c. Click OK.
   d. Double-click Load.
7. Expand Virtual User:
   a. Double-click Options and refer to Virtual Users.
   b. Click OK.
   c. Double-click Run to launch the workload.
8. Repeat Steps 1-7 for each instance. Be sure to navigate to the correct path for each HammerDB instance and change the Oracle Service Name to reflect the correct database under test.

11.2 Review the Results

When each test completes, the results will be listed in two metrics: new orders per minute (NOPM) and TPM. NOPM is a score specific to the TPROC-C benchmark of HammerDB and measures the rate at which the virtual users across the warehouses create new orders. TPM measures all transactions over a period of time on the database, regardless of the source.

To find the results, look in one of two places, the GUI or the log files. The Virtual User Output tab of the GUI will list both metrics in the Virtual User 1-MONITOR box. Look for the line TEST RESULT. The log files contain the same statements and can be found in the Datagen log path.

12 Appendix A: Create the vSAN Cluster

1. **Connect to the same PCIe root complex.** For maximum performance, ensure the intended cache and capacity disks for the same vSAN disk group connect to the same PCIe root complex. Failure to do so may result in storage traffic between cache and capacity disks crossing CPU boundaries and incurring latency penalties. Accomplish verification with a combination of visual or logical inspection.
   a. **Visual inspection**
      i. Note the physical drive placement of each drive and its serial number.
      ii. Consult the platform or server manual for PCIe layout related to CPU.
   b. **Logical verification**
      i. Establish an SSH session with the host.
      ii. Identify the Bus, Device, Function (BDF) addresses of each storage device and the Virtual Machine Host Bus Adapter (VMHBA) designation:
         1. Type "`lspci | grep vmhba`" and press Enter.
         2. The output below is captured from a dual-socket system used in the testing. Devices with a bus address less than 0x7f (decimal 127) are located on the first bus (CPU0). Devices with an address of 0x80 and above are on the second bus (CPU1). The output below demonstrates there is precisely one Intel Optane SSD and three NAND SSDs for each bus or CPU.

```
[root@fm42mns001:~] lspci | grep vmhba
0000:00:17.0 SATA controller: Intel Corporation Lewisburg SATA AHCI Controller [vmhba0]
0000:65:00.0 Non-Volatile memory controller: Intel Corporation NVMe Datacenter SSD [Optane] [vmhba1]
0000:66:00.0 Non-Volatile memory controller: Intel Corporation NVMe Datacenter SSD [3DNAND] SE 2.5" U.2 (P5510) [vmhba2]
0000:67:00.0 Non-Volatile memory controller: Intel Corporation NVMe Datacenter SSD [3DNAND] SE 2.5" U.2 (P5510) [vmhba3]
0000:68:00.0 Non-Volatile memory controller: Intel Corporation NVMe Datacenter SSD [3DNAND] SE 2.5" U.2 (P5510) [vmhba4]
0000:ca:00.0 Non-Volatile memory controller: Intel Corporation PCIe Data Center SSD [vmhba9]
0000:e3:00.0 Non-Volatile memory controller: Intel Corporation NVMe Datacenter SSD [Optane] [vmhba5]
0000:e4:00.0 Non-Volatile memory controller: Intel Corporation NVMe Datacenter SSD [3DNAND] SE 2.5" U.2 (P5510) [vmhba6]
0000:e5:00.0 Non-Volatile memory controller: Intel Corporation NVMe Datacenter SSD [3DNAND] SE 2.5" U.2 (P5510) [vmhba7]
0000:e6:00.0 Non-Volatile memory controller: Intel Corporation NVMe Datacenter SSD [3DNAND] SE 2.5" U.2 (P5510) [vmhba8]
```

      iii. Identify the storage paths for each VMHBA and device serial number:
         1. Type "`esxcli storage core path list | grep -B1 Runtime`" and press Enter.
         2. Note the model and serial numbers embedded in the UID for each VMHBA and save them for when you create the vSAN disk groups.

```
[root@fm42mns001:~] esxcli storage core path list | grep -B1 Runtime
   UID: pcie.6800-pcie.0:0-t10.NVMe____INTEL_SSDPF2KX038TZ____________________000330107BE4D25C
   Runtime Name: vmhba4:C0:T0:L0
--
   UID: pcie.e500-pcie.0:0-t10.NVMe____INTEL_SSDPF2KX038TZ____________________00032FFC13E4D25C
   Runtime Name: vmhba7:C0:T0:L0
--
   UID: pcie.6700-pcie.0:0-t10.NVMe____INTEL_SSDPF2KX038TZ____________________0003301242E4D25C
   Runtime Name: vmhba3:C0:T0:L0
--
   UID: pcie.e400-pcie.0:0-t10.NVMe____INTEL_SSDPF2KX038TZ____________________00032FFC2FE4D25C
   Runtime Name: vmhba6:C0:T0:L0
--
   UID: pcie.ca00-pcie.0:0-t10.NVMe____MO0800KEFHP____________________________PHFT640500E0800HGN__00000001
   Runtime Name: vmhba9:C0:T0:L0
--
   UID: pcie.6600-pcie.0:0-t10.NVMe____INTEL_SSDPF2KX038TZ____________________00033012C3E4D25C
   Runtime Name: vmhba2:C0:T0:L0
--
   UID: iqn.1998-01.com.vmware:fm42mns001.fm.intel.com:703875036:65-00023d000003,iqn.1991-05.com.microsoft:
   Runtime Name: vmhba65:C0:T0:L0
--
   UID: pcie.e300-pcie.0:0-t10.NVMe____INTEL_SSDPF21Q400GB____________________00026E3237E4D25C
   Runtime Name: vmhba5:C0:T0:L0
--
   UID: pcie.e600-pcie.0:0-t10.NVMe____INTEL_SSDPF2KX038TZ____________________00032FFC35E4D25C
   Runtime Name: vmhba8:C0:T0:L0
--
   UID: pcie.6500-pcie.0:0-t10.NVMe____INTEL_SSDPF21Q400GB____________________00026DEB1CE4D25C
   Runtime Name: vmhba1:C0:T0:L0
```

2. **Create a new cluster**
    a. Accept the default settings (all options disabled or unchecked).
    b. Add target hosts to the cluster.

3. **Enable vSAN**
    a. Do not enable any data services (for example, compression, deduplication and encryption).

4. **Create vSAN disk groups**
    a. Create disk groups one at a time for each host.
    b. Consult the list from the previous steps to select cache and capacity drives on the same PCIe bus.
    c. Repeat until each host has two disk groups.
    d. This can also be verified after creation in "Disk Management" for the cluster:
        i. Select the cluster and navigate to "Configure > vSAN > Disk Management".
        ii. Select a host and click the "VIEW DISKS" link.

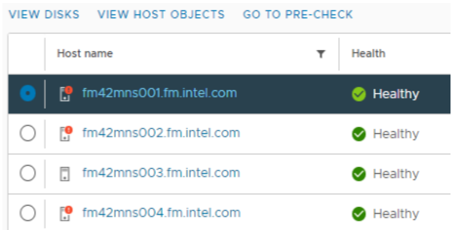

        iii. Expand one of the disk groups by clicking the arrow on the left.

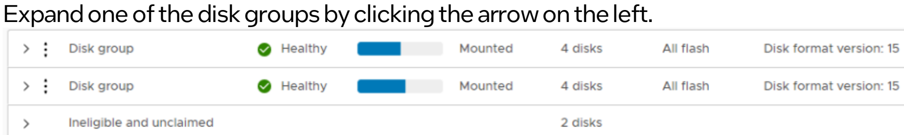

        iv. Expand the Name column to reveal the UID of each device and note the serial numbers.

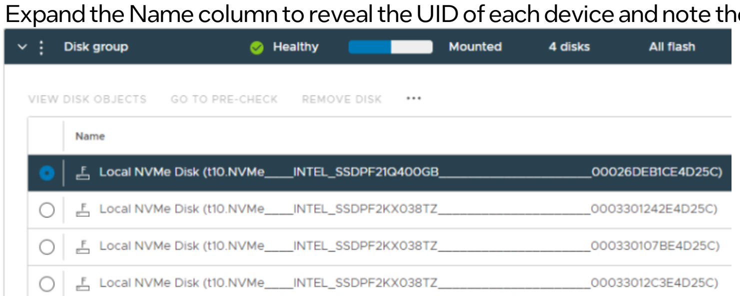

5. **Verify that the vSAN Default Storage Policy matches the configuration in the previous section.**

Revision History

Document Number	Revision Number	Description	Date
	1.0	First Release	June 2022