Best Practice

Workload Optimization _ Best Practice for the QxStack VMware vSAN-based Solution

Table of Contents

Tables

Figures

1 Executive Summary

In a highly competitive business environment, many enterprises rely heavily on business-critical applications (BCA) for resource planning and decision making. However, enterprises often encounter issues such as storage inefficiency, complex management, inability to scale-up, and high OPEX and CAPEX when running BCAs, prompting them to consider running their business-critical applications in a virtualized environment.

Of all business-critical applications, data collection and data analysis play significant roles in business management and strategy forming. More and more companies have adopted database management systems like Oracle to manage their data. Oracle is the most popular database management system among enterprises to date, offering flexible deployment, on-demand scalability, and high-performance database solutions. However, establishing and maintaining data storage is not an easy task. The efficiency of database management systems must be enhanced to boost datacenter performance.

To ensure customers enjoy optimized and promised performance in database systems, Quanta Cloud Technology (QCT), a global data center solution provider, has integrated its pre-validated hardware with VMware vSAN™ to provide our customers a high-performance and cost-efficient database best practice for data management. With a Hyper-Converged Infrastructure (HCI), QCT's best practice allows administrators and engineers to deploy and run their database workload in a way that is tailored to their needs. Besides, with QCT's unique tuning approach, we boost the performance of databases to the extreme, making it more powerful to fulfill your database workload.

With aforementioned benefits and outstanding performance, QCT's best practice is a valid approach to help customers fulfill database workload more efficiently.

2 Overview of Best Practice

The purpose of this document is to introduce the best practice of QxStack VMware vSAN-based solution with QCT's tuning approach for workload optimization. This best practice will illustrate how QCT provides you an extremely high-performance software-defined data center (SDDC) solution with validated configuration and parameter-tuning skills.

2.1 Scope

This best practice

- illustrates why QCT adopts the architecture of vSAN HCI to run business-critical applications,
- introduces the techniques of QCT workload optimization,
- illustrates the hardware configuration and software components discreetly selected by QCT to provide a cost-efficient vSAN-based solution, and
- demonstrates the validation of QCT's tuning capabilities and show the test results of best practice.

2.2 Audience

The intended audience of this document are IT professionals, technical architects, and sales engineers. This document is to assist them in planning, architecting and implementing business-critical applications with vSAN.

3 Solution Overview

This section introduces the QxStack VMware vSAN-based solution. Specifically, it demonstrates the benefits of QCT's architecture design as well as skilled tuning approach.

3.1 QxStack VMware vSAN-based solution

QCT leverages its well-designed 2nd Generation server platform and the market leading virtualization software developed by VMware to deliver an extremely high-performance vSAN-based solution, providing a reliable choice to our customers (see Fig. 1).

This solution addresses the common issues enterprises face for business-critical applications such as storage inefficiency, complex management and inability to scale-up and scale-out. Furthermore, with skilled tuning approach, QCT can boost the performance of workloads to the extreme, making it more powerful to fulfill customers' business demands.

The following sections provide a solution overview, highlighting some benefits that customers can enjoy while solving their problems.

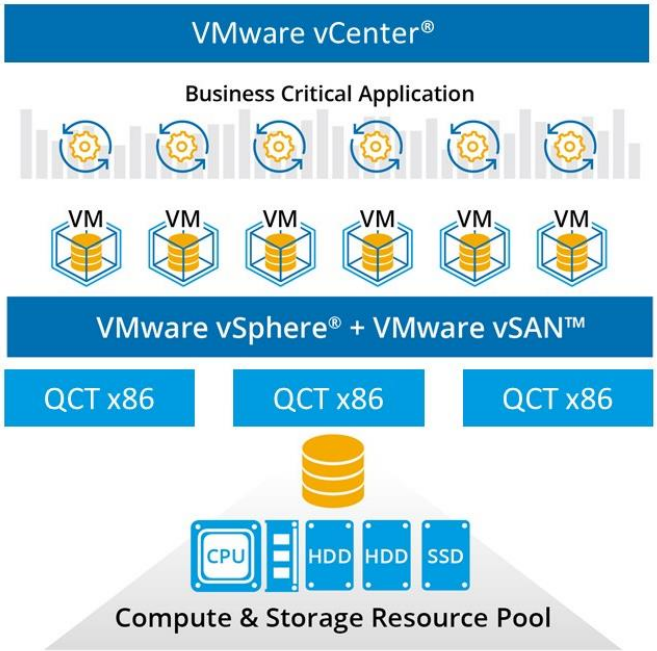

Figure 1. Solution Architecture.

3.1.1 Reliability - vSAN™ Certification

QCT has validated this solution to ensure its reliability, allowing our customers to focus on strategic and productive tasks.

Table 1. QxStack vSAN™- pre-validated components.

Components	Details	QTY (Total)	QTY (Per node)
System	Model: QuantaGrid D52B-1U System Type: Rackmount	3	1
CPU	Intel® Xeon® Gold 6252 CPU @ 2.10 GHz / 48C/96T	6	2
Memory	32GB 2666MHz DDR4 RDIMM	54	18
Caching Tier	Model : Intel® SSD DC P4800X Series SSDPED1K375GA (375 GB, AIC) Partner Name: Intel Device Type: NVMe Capacity: 375 GB Performance Class: Class F: 100,000+ writes per second TBW Endurance Class: Endurance Class D >=7300 TBW	6	2
Capacity Tier	Model : Intel® SSD D3-S4510 Series SSDSC2KB038T8 (3.84TB, 2.5”) Partner Name: Intel Device Type: SATA Capacity: 3840 GB Performance Class: Class D: 20,000-30,000 writes per second TBW Endurance Class: Endurance Class D >=7300 TBW	18	6
NIC 1	Model: Intel(R) Ethernet Controller XXV710 for 25GbE SFP28	3	1
NIC 2	Model: Intel(R) Ethernet Connection X722 for 10GbE SFP+	3	1
Boot Device	SATADOM 32GB	3	1

3.1.2 Scalability, Manageability and High Availability- HCI

Traditionally, IT technicians face the challenges of resource management and scalability since compute and storage resources are separated. Hyper-Converged Infrastructure (HCI) is a novel technology that can integrate compute, storage, and virtualization resources in a single hardware box as shown in Fig. 2. Every single node is capable of delivering compute and storage resources at the same time.

Furthermore, when enterprises run business critical applications, they must take availability, recoverability, scalability and security issues into account. By adopting a Hyper-Converged Infrastructure, users can address these issues easily with simplified operations and at a lower TCO. Key benefits are listed below:

Scalability

The hyper-converged infrastructure integrates compute and storage resources into a basic unit, which is also referred to as a building block. By implementing a clustered architecture, users can add more building blocks to the cluster to expand the overall performance and capacity.

Manageability

In legacy converged architectures, compute and storage resources are provided separately by servers and storage devices, which have independent configuration settings and operation management tools. That means IT administrators must manage two sets of devices using different management tools. In a hyper-converged infrastructure, full “policy-driven management” has been realized to ease the burden of IT administrators, who now only need to define their compute and storage resources. The allocation process can be automatically completed through a single management portal, which significantly reduces the management efforts.

VMware vSAN also features enhanced functions of snapshots, backup, recovery and cloning, making day-2 operation of business-critical applications more convenient.

High Availability

Ensuring an agreed level of operational performance is essential when running business critical applications. A VMware vSAN hyper-converged infrastructure automatically takes resynchronization actions to re-establish the full compliance of the storage protection policies assigned to the affected VM. In the meantime, resynchronization and guest VM traffic are maintained at a balance to ensure the servers can deliver satisfactory levels of performance at all times.

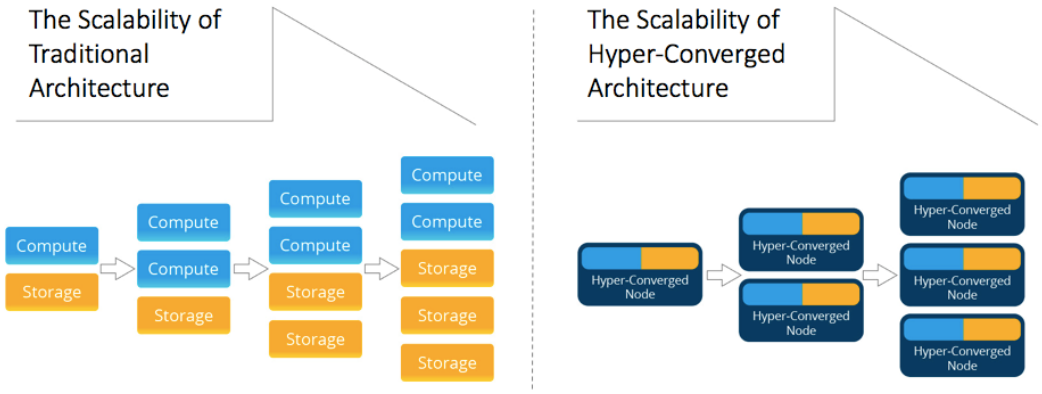

Figure 2. Comparison between a traditional architecture and HCI.

3.1.3 Ultra-performance with Cost-efficiency

The QxStack VMware vSAN-based solution adopts high-density, high-performance, energy-efficient QCT servers powered by Intel® Xeon® Scalable processors as the platform. Compared with an all NVMe vSAN-based solution, the solution that adopts a combination of NVMe and SATA SSD is more cost-efficient.

3.2 QCT Tuning Approach

Workload optimization refers to leveraging the underlying hardware, software, middleware layers to achieve higher performance. QCT has rich experiences in performance tuning as we have helped numerous customers elevating their system performance. Below are some key steps to consider when running the workload tuning process:

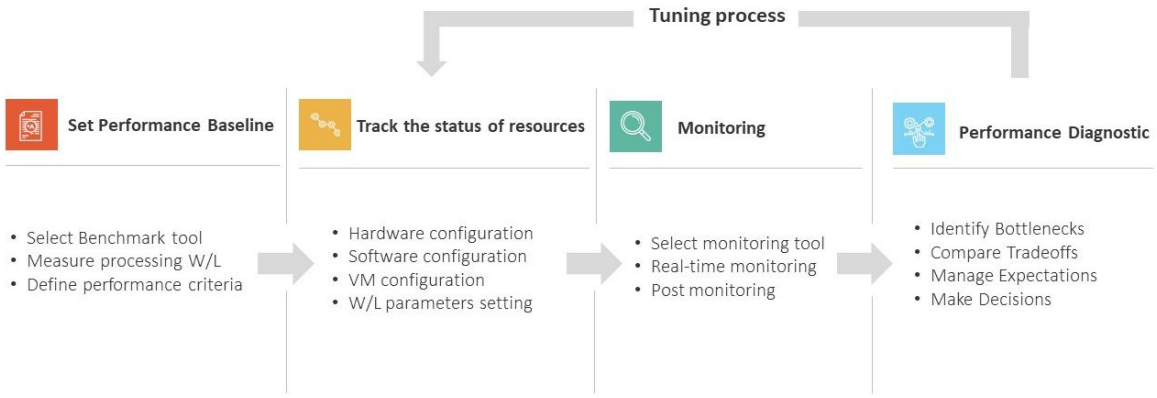

Figure 3. QCT workload tuning process

Set Performance Baseline

It is essential to set a performance baseline when tuning a database, since it is common for DBAs to find that an optimization technique has little, if not negative effect, on the workload. It is easier to determine the bottlenecks or the effectiveness of the parameters by comparing current performance and the performance baseline. It can also serve as a reference when DBAs try to manage and monitor database systems.

Track the Status of Resources

While performance baselines record the statistics about the operating status of servers and databases, providing the information to clarify the current or potential issues is also crucial to tracking resource usage. From QCT's experience, it is essential to track the status of resources from multiple aspects:

For **hardware configuration**, QCT will track the basic hardware resource usage such as power setting and the status of host, network and CPU. After the adoption of the VMware vSAN hyper-converged infrastructure (HCI), the usage of storage resource must also be tracked, including the disk type (NVMe, SAS or SATA), the quantity of disk groups per host, or the quantity of capacity disks per disk group.

For **software configuration**, it is important to track the status of software resources and setting. For example, space saving, data security or other settings including stripe width and the IOPS limit for object.

For **VM configuration**, or the arrangement of the resources including allocated processors, memory, network, disk and OS assigned to the virtual machine, here are also some VM resources needed to track: VM sizing, vCPU and vRAM usage, and virtual disk setting.

As for **workload parameter setting**, in addition to the parameter setting of the Oracle database itself, some other parameters from different aspects also need to be considered, for example, BIOS, ESXi, vSAN, and the benchmark tool.

Monitoring

After setting the performance baseline and tracking the resource status, the next step would be resource monitoring, which becomes increasingly important as it is the easiest way to evaluate performance, identify potential bottlenecks, and come up with a solution for troubleshooting and fine-tuning. Numerous monitoring tools that are currently available can be divided into real-time monitoring and post-monitoring ones. Making good use of these tools is essential throughout the workload optimization process.

Performance Diagnostics

Resource monitoring makes bottleneck identification easier. Any bottleneck that may undermine the performance must be identified and eliminated to minimize their impact on the workload. After identifying the bottlenecks, parameter tuning would be the next step. Bear in mind that when choosing which parameters should be tuned, you may have to consider trade-offs.

Based on the rich experience helping customers boost performance to fulfill their business demands, QCT has developed a unique tuning approach and established a list of the best-fit parameters by workloads and by use case. For more detailed information about tuning items, please refer to Chapter 5.

4 Solution Architecture

This session introduces the solution architecture, including hardware architecture and software configuration.

4.1 Hardware Architecture

4.1.1 Compute: Intel® Xeon® Scalable Processor

Tailored for a software-defined storage solution, QuantaGrid D52B-1U features ultimate compute and storage density in a 1U platform which redefines the physical limitation. QuantaGrid D52B-1U is a rackmount server with Intel® Xeon® Processor Scalable Family, delivering fast socket interconnect, 1.5x memory bandwidth, and 2x Floating-Point Operations per second (FLOPs) peak performance. Intel® Xeon® Scalable processors can deliver up to 112 vCPUs per server and 3.9x higher virtualized throughput compared to the previous platforms based on the Intel® Xeon® processor E5. The compute capability can empower VMware® vSAN™ by supporting a wide range of critical workloads with low latency in the modern data center.

4.1.2 Storage: Intel® Optane™ NVMe SSD

Caching is prerequisite not only in today's software-defined storage solution but also in virtualization applications to boost performance. Enterprises that value performance can benefit from the cache tier using the Intel® SSD Data Center (DC) Family with Non-Volatile Memory Express (NVMe). Adopting Optane™ SSD DC P4800X as the cache tier in vSAN™ can deliver an extremely high-performance and ultra-low latency while running the write-intensive workload.

1. Intel® Optane™ SSD DC P4800X combines the attributes of both memory and storage, significantly reducing the latency of sensitive workload and accelerating overall deployment.
2. Intel® Optane™ SSD DC P4800X is able to deliver 5-8x faster performance to low queue depth workload, exhibiting exceedingly high throughput.

4.1.3 Network: Intel® XXV710-DA2 25GbE

According to the best practice of VMware® vSAN™ configuration, at least 10 gigabit

Ethernet (GbE) is recommended on all flash vSAN™ configurations. Discreetly selected to fulfill the bandwidth demand between the server for vSAN™ traffic, Intel® Ethernet Converged Network Adapter XXV710-DA2 25GbE not only optimizes the performance but also improves the reliability for data center networking.

4.2 Software Configuration

4.2.1 VMware vSphere®

VMware vSphere® is an industry-leading virtualization platform which virtualizes and aggregates physical hardware to provide a virtual resource pool to the data center. By leveraging the virtualization technology, vSphere® can provide a highly available, efficient, centralized infrastructure for IT administrators to deliver flexible, reliable services. VMware vSphere® provides the following benefits:

1. Using proactive high-availability technology to prevent machine downtime.
2. Using predictive load balancing technology to fully exert data center resources.
3. Simplifying user experience through delivering large-scale automation and management.
4. Leveraging virtual machine level encryption technology to reduce security risk.
5. Using RESTful API to promote IT automation and business flexibility.

4.2.2 VMware vCenter Server®

vCenter server® is a platform which provides centralized management and operation for VMware® virtualization environment. By aggregating all the virtual resources, vCenter server® can provide resource provisioning and monitoring. It is capable of provisioning compute, storage, and other resources to the virtual machine and enabling High Availability (HA), Distributed Resource Scheduler (DRS), vMotion, etc. Currently, in vCenter® 7.0.

4.2.3 VMware ESXi™

VMware® ESXi™ is an industry-leading hypervisor installed on a bare-metal physical server. ESXi™ has its own kernel, called VMkernal, based on Linux kernel. ESXi™ enables a virtualization technology that is different from the traditional hardware architecture, which includes compute, storage, and networking. By sharing the resources of a single hardware unit across multiple environments, a physical system (x86 server) is capable of executing multiple virtual machines with different operation systems (OS). VMware® proposed that “virtualization is the process of creating a

software-based (or virtual) representation of something rather than a physical one." Virtualization can be applied to applications, servers, storage, and networks. It is a single effective way to reduce IT expenses and boost efficiency and agility for all-sized businesses.

4.2.4 VMware vSAN™

vSAN™ is a software-defined storage solution built in the vSphere® kernel. Tightly integrated with the hypervisor to minimize the CPU and memory overhead, it optimizes the data I/O path to deliver an outstanding performance. vSAN™ is a hyper-converged solution specifically designed for virtual machines. It minimizes the effort to configure the storage and simplifies virtual machine deployment.

vSAN™ is built on the industry-standard x86 server and leverages local storage on the server. By using the virtualization technology, the under-layer physical disk is transformed to abstract resource pool, aggregated into a virtual resource for providing a shared storage to the server in the cluster. It also makes non-disruptive expansion capacity possible by adding hosts to a cluster or adding disks to a host. SSDs are used as a cache tier to accelerate the I/O performance and HDDs as capacity tier to store data. By leveraging the flash device on server side, vSAN™ can accelerate the read/write I/O processing speed and minimize the storage latency. Administrators can use vSAN™ to define the requirements of virtual machine storage such as performance and availability. The policy requirements are delivered to the vSAN™ layer through Storage Policy-Based Management (SPBM). Hence, SPBM plays an important role to bridge upper-layer applications and under-layer storage devices. Through pre-defined storage policy, SPBM will drive the vSAN™ layer to adopt these policies when users provision the virtual machine.

5 Solution Validation

In this section, the testing process and test results are presented to illustrate the performance and reliability of this best practice.

5.1 Test Overview

In this test, an Oracle database was adopted to demonstrate the business-critical application in a business environment. Oracle is a database commonly used for running online transaction processing (OLTP) database workloads. An OLTP workload refers to a workload that receives both requests for data and changes to these data (or transactions) from multiple users over time. Oracle is the most widely used RDBMS (Relational Database Management System). The purpose of this test was to determine the maximum performance of Oracle Database instances when running the OLTP workload. The Oracle Database instances were hosted by one or more virtual machines. HammerDB served as a benchmark tool to generate workload and measure performance during the test. This benchmark tool was used to test an OLTP database system with performance metrics such as TPM (Transaction Per Minute) and NOPM (New Order Per Minute). HammerDB adopted a simplified version of **TPC-C industry standard benchmark** to simulate OLTP workload.

As shown in Figs. 4-5, TPC-C simulates a computer system that helps a wholesaler deliver goods. It assumes there is a large commodity wholesaler that has several warehouses in different regions. Each warehouse is responsible for supplying 10 sales points, while each sales point serves 3000 customers. However, a few items in a customer's order are not in stock in the local warehouse, and need to be supplied by a warehouse in another region. TPC-C benchmark tracks five stages of transaction behaviors: new order, payment, order Status, delivery and stock-level. These transaction behaviors combined simulate a complet and complex cycle of real-world business operations.

During the test, QCT focused on finding the best performance value for Oracle database. By tuning the parameters of BIOS, vSphere, vSAN, OS, HammerDB and Oracle database, QCT boosted the performance to its maximum. Detailed information is provided in the next session.

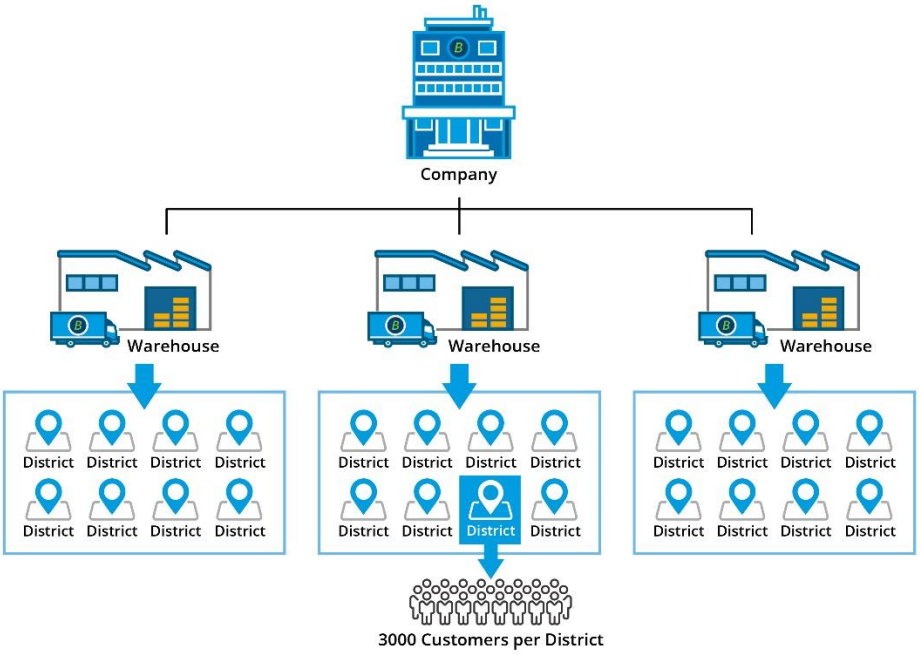

Figure 4. TPC-C Company Structure

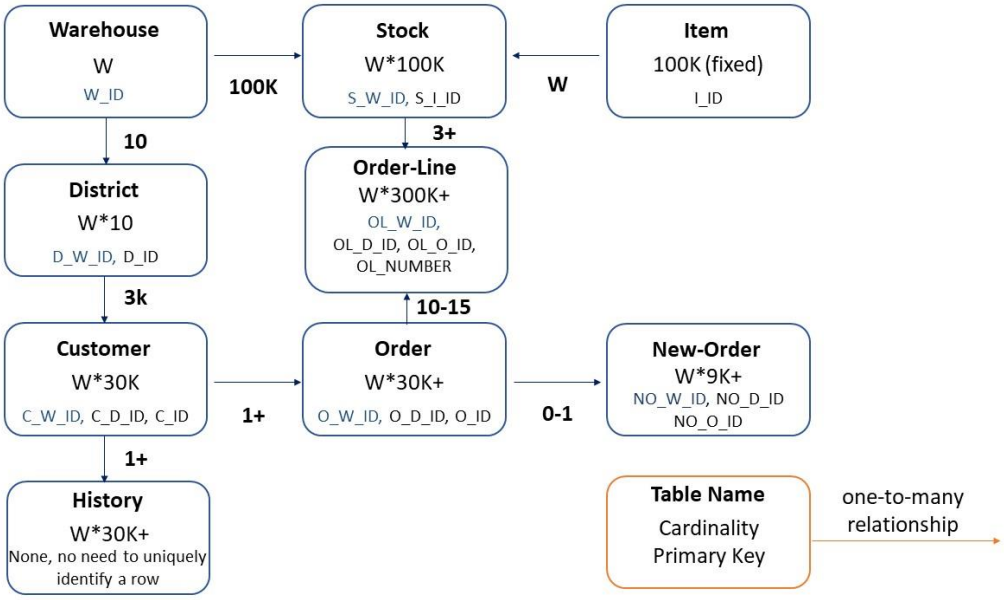

Figure 5. TPC-C Database Schema

5.2 Test Configuration

Figure 6 shows a three-node cluster configuration is used to form a VMware vSAN cluster during the test. The vCenter appliance is installed on a separate infrastructure. Six HammerDB clients that generate workloads towards Oracle Databases are grouped in separated infrastructures as well. One of the vSAN cluster physical hosts is selected as the target of this test.

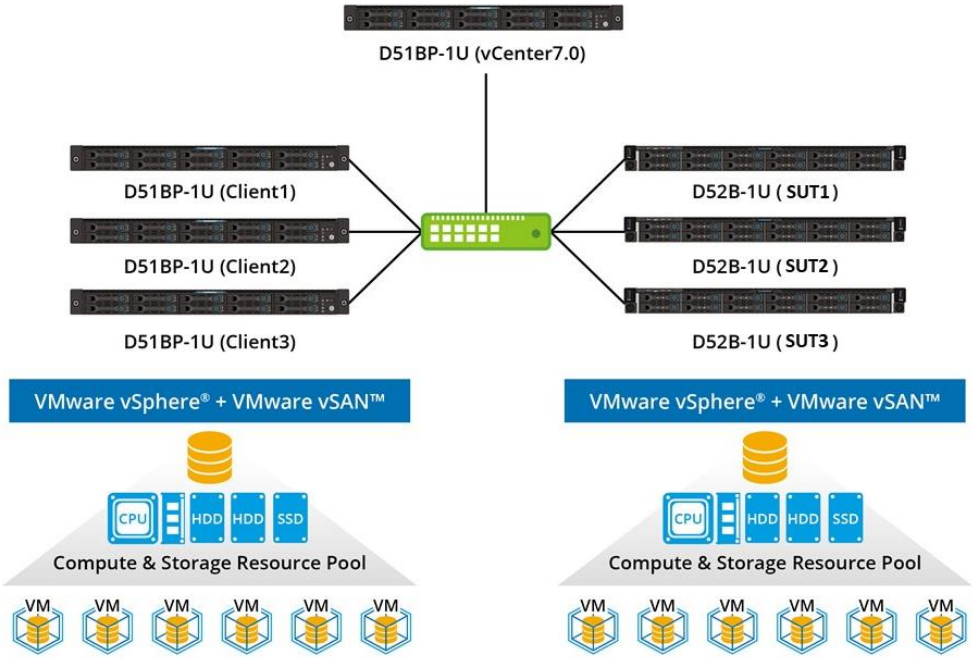

Figure 6. Cluster topology

Besides, a well-designed network topology ensures that vSAN™ traffic can run efficiently and correctly between the hosts to deliver high availability. VMware® suggests to adopt a 10 Gigabit Ethernet (GbE) network for vSAN™. In QCT's network design, NIC teaming is adopted to team the two NICs together, thereby avoiding bandwidth bottlenecks for the vSAN service running on different switches and enhance availability. Detailed network topology is shown in Fig.7.

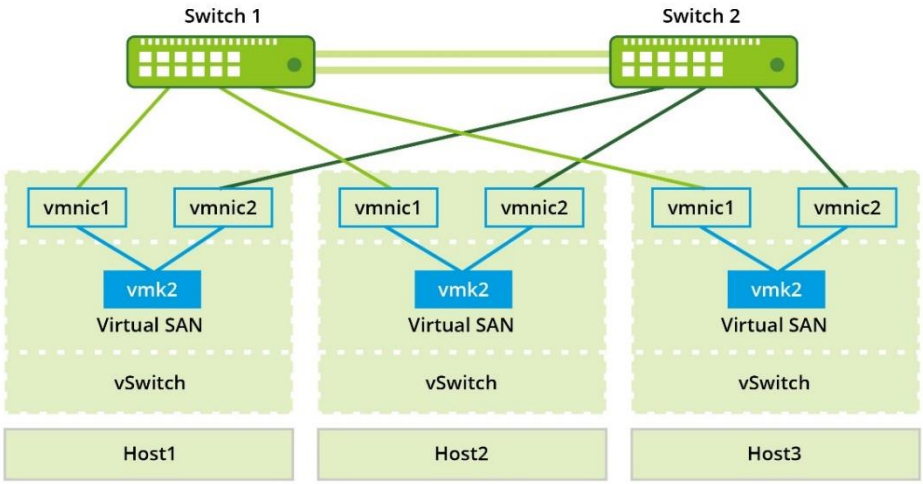

Figure 7. vSAN Network Topology

5.2.1 Hardware Configuration

Table 2. Bill of Material for SUT (D52B-1U)

Component	Description	QTY (Total)	QTY (per node)
System	QuantaGrid D52B-1U	3	1
CPU	Intel(R) Xeon(R) Gold 6252 CPU @ 2.10GHz	6	2
DIMM	RDIMM DDR4 32GB 2666MHZ	54	18
Controller	QS-3216	3	1
SSD cache	Intel NVMe Datacenter SSD [Optane] x4 AIC (P4800X) 375GB	6	2
SSD data	Intel SSD DC S4510 3.84TB	18	6
Boot device	SATADOM 32GB	3	1
NIC 1	Intel(R) Ethernet Controller XXV710 for 25GbE SFP28	3	1
NIC 2	Intel(R) Ethernet Connection X722 for 10GbE SFP+	3	1

Table 3. Bill of Material for Client (D51BP-1U)

Component	Description	QTY (Total)	QTY (per node)
System	QuantaGrid D51BP-1U	3	1
CPU	Intel(R) Xeon(R) CPU E5-2699 v4 @ 2.20GHz	6	2
DIMM	RDIMM DDR4 32GB 2133MHZ	18	6
SSD cache	Intel NVMe SSD DC P3700 SSD [2.5" SFF] 800GB	6	2
SSD data	Intel NVMe SSD DC P3500 SSD [2.5" SFF] 2TB	18	6
Boot device	SATADOM 32GB	3	1
NIC	Intel(R) 82599 10 Gigabit Dual Port Network Connection	3	1

5.2.2 Software Configuration

Table 4. Software Versions

Component	Versions
VMware	
VMware vSphere 7.0	7.0
VMware vCenter Server Appliance 7.0	7.0
VMware vSAN 7.0	7.0
VMware vDS 7.0.0	7.0.0
HammerDB	
HammerDB v3.3	6

Table 5. Guest OS Configuration

Number of HammerDB clients per VM	Number of hosts	Number of VMs	vCPUs	vRAM	Oracle software
1	3	6	48	256GB	Oracle GRID 12.2.0.1.0
					Oracle Database 12c Release 2 Enterprise Edition 12.2.0.1.0
					Oracle Linux 7.4

5.3 Test Cases

Based on the guest OS configuration (see Table 5), QCT defined two VM configurations according to different HammerDB settings shown in Table 6.

To determine the maximum performance for the best practice, QCT developed three test cases—**Baseline, Tuned and Best result**—based on the two VM configurations. In Table 7, you can find more details about the three test cases.

Table 6. HammerDB configurations

VM ID	TPCC database instances per VM	Number of warehouses per VM	Number of active users per VM	Number of hosts	Number of VMs
1	1	50	250	3	6
2	1	400	100	3	6

Table 7. Test cases

Test case	Description
Baseline	Based on VM ID1, with all parameters set to the default setting.
Tuned	Based on VM ID1, with VMware vSAN, vSphere, HammerDB, BIOS, Network, Oracle parameter tuned.
Best result	Based on VM ID2, find the best TPM and NOPM without any restriction.

Below are the detailed tuning items we've recorded from the three test cases:

Table 8. Parameter tuning

Parameter	Baseline	Tuned	Best result
VMware vSAN			
FTT (Failures to Tolerate)	1	0	0
vSAN Striping	1	12	12
vSAN block size	32KB	32KB	32KB
Object space reservation	Thin provisioning	Thick provisioning	Thick provisioning
Disk Groups	2	2	2

HammerDB 3.3			
Timed Driver Script	Enable	Enable	Enable
Checkpoint when complete	Enable	Enable	Enable
Use All Warehouse	Enable	Enable	Enable
Minutes of Rampup Time	3	3	2
Minutes for Test Duration	5	5	5
Partition Order Line Table	N/A	N/A	Enable
Order Line Tablespace	N/A	N/A	Orders
BIOS Setting			
Pwe and Perf Profile	Custom	High Performance	High Performance
NUMA	Enabled	Disabled	Disabled
Hardware PM state Control > Hardware P-State	Out of Band Mode	Native Mode	Native Mode
Network			
Jumbo Frame (MTU)	1500	9000	9000
vDistributed Switch for VM Network	N/A	2	2
vSphere			
Power management	Balanced	High performance	High performance
Oracle Database VM			
Cores per socket	24	24	24
vCPU	48	48	48
CPU pinning VM 1,3,5	N/A	0-47	0-47
CPU pinning VM 2,4,6	N/A	48-95	48-95
Memory	256GB	256GB	256GB
Memory Reservation	N/A	Enable	Enable
Filesystem used for Oracle disks	ASM	ASM	ASM
VMDK / system disk	300GB	300GB	300GB
VMDK / Oracle Software Disk (u01 folder)	150GB	150GB	150GB
VMDK / ASM disk	350GB*1	350GB*8	350GB*8
+DATA ASM disk group	150GB*1	150GB*8	150GB*8
+FRA ASM disk group	130GB*1	130GB*8	130GB*8
+CRS ASM disk group	70GB*1	70GB*8	70GB*8
Filesystem used for / disks	XFS	XFS	XFS
Oracle ASM disk block size	4MB	1MB	1MB
Oracle Database block size	8KB	8KB	8KB
Oracle Global Parameter Tuning			

Check DB block content
Enlarge cache size
Increase writer processes
Enlarge the size of the in-memory Active Session History buffers
Increase the disk sector size
Change logging behavior
Enable system auditing
Stop auto tuning
Max cursor per session
Change redo log & checkpoint setting
Increase DDL waiting time but lower the DML lock quantity
Lock entire SGA in physical memory (cannot swap data into hard drive)
Max. number of concurrent active transactions
Reduce the undo retention time.
Add log write waiting time to sync the log
Use more memory in order to get faster execution
Disable system auditing
Decrease the number of active transactions per rollback segment
Enable in memory tracing
Enable reliable latch waits
Don't track dependency for Replication parallel propagation
Disable Advanced Queue
Enable system auditing
Disables the tracking of unrecoverable (NOLOGGING) direct-path insert and load operations.
Increase the user processes
Enable both asynchronous and direct I/O on file system files.
Create segment without time deferred
Automate the parallelism degree
Increase dynamic sampling
Disable the query optimizer

5.4 Test Results

Based on the aforementioned test configurations and tuning approach, QCT validated the capability of workload optimization using the three test cases—**Baseline**, **Tuned**, and **Best Result**—as shown in Table 7. Through tuning the parameters of VMware vSAN, vSphere, HammerDB, BIOS, Network and Oracle, the best result test case achieved 7.71 times higher TPM and 7.67 times NOPM when compared to the baseline case (see Fig. 8). This test result demonstrates that QCT has come up with the best practice to boost the performance to its extreme, exerting the best use of the hardware platform.

Table 9. Oracle Test Results

Test case	Warehouse	User	TPM	NOPM
Baseline	50	250	1,365,703	463,771
Tuned	50	250	9,922,889	3,337,966
Best Result	400	100	10,542,613	3,559,075

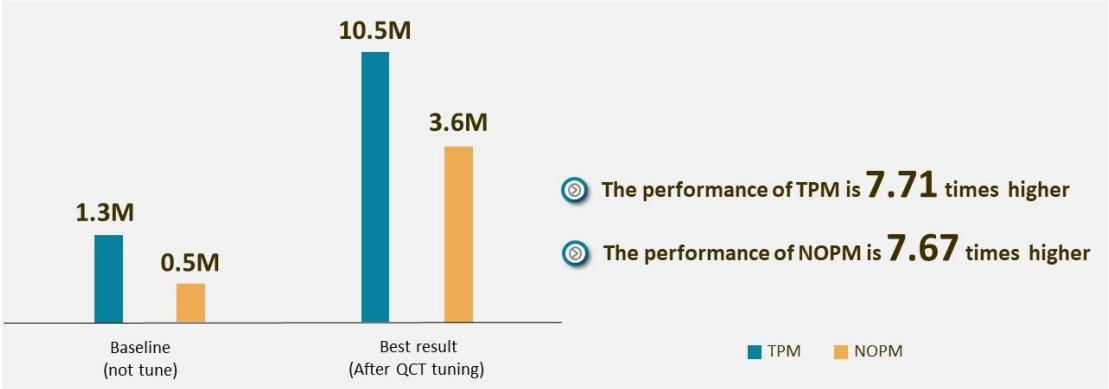

Figure 8. Testing Performance

6 Conclusion

In this age of information explosion, many enterprises have considered running their business-critical applications in a virtualization environment for easier management, better scalability as well as a lower TCO. QCT can offer you a flexible and innovative solution based on the best practice to achieve the aforementioned benefits.

QCT has validated that **QxStack VMware vSAN-based solution** can deliver its best practice with a brilliant architecture with vSAN HCI, supporting the operation of business-critical applications. With a reliable hardware platform and skilled tuning experience and approach, QCT can help its customers to achieve ultimate performance with the benefit of cost-efficiency. QCT's best practice will definitely be your first choice when running business-critical applications.

QCT appreciates your feedback on our way of continued innovation. For further inquiry, please visit
https://go.qct.io/solutions/enterprise-private-cloud/qxstack-vmware-edition-vsan-readynode/

7 References

1. **Oracle Database on VMware vSAN 6.7**

   https://storagehub.vmware.com/t/oracle-database-on-vmware-vsan-6-7/

2. **Oracle on VMware vSAN solution overview**

   https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/products/vsan/vmw-solution-overview-oracle-on-vsan.pdf

3. **Database Workload Optimization**

   https://link.springer.com/chapter/10.1007/978-1-4302-1903-3_15

4. **SQL Server performance tuning best practices for DBAs**
   https://searchsqlserver.techtarget.com/tip/SQL-Server-performance-tuning-best-practices-for-DBAs

5. **Oracle Database monitoring, tuning deliver performance goods**

   https://searchoracle.techtarget.com/ehandbook/Oracle-Database-monitoring-tuning-deliver-performance-goods

6. **Reasons to move business-critical applications to virtual machines**

   https://datadepositbox.com/reasons-to-move-business-critical-applications-to-virtual-machines/

7. **HammerDB database testing tool**
   http://www.hammerdb.com/about.html

8. **Tegile Drives Oracle @ the Speed of Memory**
   https://stratsolutions.com/tegile-drives-oracle-the-speed-of-memory/

9. **Intel® SSD Data Center Family**
   https://www-ssl.intel.com/content/www/us/en/solid-state-drives/data-center-family.html

10. **QuantaGrid D52B-1U**
    https://www.qct.io/product/index/Server/rackmount-server/1U-Rackmount-Server/QuantaGrid-D52B-1U

11. **Automatic Storage Management Administrator's Guide**
https://docs.oracle.com/database/121/OSTMG/GUID-BC6544D7-6D59-42B3-AE1F-4201D3459ADD.htm#OSTMG10042

12. **Oracle 12cR2 Installation on Linux with ASM**
https://support.dbagenesis.com/knowledge-base/oracle-12c-installation-on-linux-with-asm/#install-oracle-software

13. **Purge ARCHIVELOG automatically**
https://dba.stackexchange.com/questions/103586/purge-archivelog-automatically

14. **Accelerate Oracle Database with the Next-Gen FlashArray//X**
https://blog.purestorage.com/oracle-database-flash-array-xr3/

15. **Tegile's Oracle TPC-H test report**
https://blog.westerndigital.com/oracle-analytical-processing-nvme-test-results-guide/

16. **Netapp's Oracle benchmark test report (multiple benchmark tool)**
https://www.netapp.com/us/media/tr-4794.pdf

17. **oracle add redo log group member**
https://blogsendak.pixnet.net/blog/post/90784438

18. **How to Install Oracle Database 12c on RHEL/CentOS 7**
https://www.tecmint.com/install-oracle-database-12c-on-centos-7/

19. **Driving hammerdbcli from a bash script**
https://www.hammerdb.com/blog/2018/10/

20. **Change redo log file size**
http://erichuang89.blogspot.com/2008/12/change-logfile-size.html

21. **oracle add redo log group member**
https://blogsendak.pixnet.net/blog/post/90784438

22. _use_adaptive_log_file_sync know more
https://www.askmaclean.com/archives/_use_adaptive_log_file_sync-know-more.html

23. Oracle® Database Administrator's Guide 11g Release 2 (11.2)
https://docs.oracle.com/html/E25494_01/restart001.htm

LEGAL DISCLAIMER

ABOUT QCT

Quanta Cloud Technology (QCT) is a global data center solution provider. We combine the efficiency of hyperscale hardware with infrastructure software from a diversity of industry leaders to solve next-generation data center design and operation challenges. QCT serves cloud service providers, telecoms and enterprises running public, hybrid and private clouds.

Product lines include hyperconverged and software-defined data center solutions as well as servers, storage, switches and integrated racks with a diverse ecosystem of hardware component and software partners. QCT designs, manufactures, integrates and services cutting-edge offerings via its own global network. The parent of QCT is Quanta Computer, Inc., a Fortune Global 500 corporation. www.QCT.io

UNITED STATES

QCT LLC., Silicon Valley office
1010 Rincon Circle, San Jose, CA 95131
TOLL-FREE: 1-855-QCT-MUST
TEL: +1-510-270-6111
FAX: +1-510-270-6161
Support: +1-510-270-6216

QCT LLC., Seattle office
13810 SE Eastgate Way, Suite 190, Building 1,
Bellevue, WA 98005
TEL: +1-425-633-1620
FAX: +1-425-633-1621

CHINA

云达科技，北京办公室（Quanta Cloud Technology）
北京市朝阳区东大桥路 12 号润诚中心 2 号楼
TEL +86-10-5920-7600
FAX +86-10-5981-7958

云达科技，杭州办公室（Quanta Cloud Technology）
浙江省杭州市西湖区古墩路浙商财富中心 4 号楼 303 室
TEL +86-571-2819-8650

JAPAN

Quanta Cloud Technology Japan 株式会社
東京都港区芝大門 2-5-8 芝大門牧田ビル 3F, 105-0012
TEL +81-3-5777-0818
FAX +81-3-5777-0819

GERMANY

Quanta Cloud Technology Germany GmbH
Hamborner Str. 55, 40472 Düsseldorf
TEL +492405-4083-1

TAIWAN

雲達科技（Quanta Cloud Technology）
桃園市龜山區文化二路 211 號 1 樓
1F, No. 211 Wenhua 2nd Rd., Guishan Dist., Taoyuan City 33377, Taiwan
TEL +886-3-286-0707
FAX +886-3-327-0001