

Hitachi Solution for Databases Using Oracle 19c Database with Hitachi Virtual Storage Platform 5500 and Hitachi Advanced Server DS7000 Series

Reference Architecture Guide

By Hitachi Vantara

October 2020

Feedback

Hitachi Vantara welcomes your feedback. Please share your thoughts by sending an email message to SolutionLab@hitachivantara.com. To assist the routing of this message, use the paper number in the subject and the title of this white paper in the text.

Revision History

Revision	Changes	Date
MK-SL-203-00	Initial release	July 16, 2020
MK-SL-203-01	Updated with AEP results	October 28, 2020

Table of Contents

Solution Overview	2
Business Benefits	2
High Level Infrastructure	2
Key Solution Components	4
Hitachi Virtual Storage Platform 5000 Series	7
Hitachi Storage Virtualization Operating System RF	7
Hitachi Ops Center	7
Hitachi Advanced Server DS7000 Series	7
Hitachi Advanced Server DS120	8
Oracle Linux	8
Oracle Database	8
VMware vSphere	8
Intel Optane Persistent Memory	9
Brocade Fibre Channel Switches from Broadcom	9
Cisco Nexus Data Center Switches	9
Solution Design	10
Storage Architecture	10
Server and Application Architecture	18
SAN Architecture	18
Network Architecture	19
Physical Network Configuration	20
Engineering Validation	22
Test Methodology	22
Database Configuration	22
Test Environment	23
Test Results	23
Conclusion	25

Hitachi Solution for Databases Using Oracle 19c Database with Hitachi Virtual Storage Platform 5500 and Hitachi Advanced Server DS7000 Series

Reference Architecture Guide

Use this reference architecture guide to see how Hitachi Solution for Databases provides a high performance, integrated, converged solution for Oracle databases using Hitachi Virtual Storage Platform 5500 (VSP 5500). This environment uses the following:

- Virtual Storage Platform 5500
- Hitachi Advanced Server DS7000 with second-generation Intel Xeon scalable processors
- Hitachi Advanced Server DS120 with Intel Xeon scalable processors.

With these products, design your infrastructure for Oracle databases to meet your requirements and budget.

This solution uses the high-performing Virtual Storage Platform 5500 with NVMe technology to boost performance. A dedicated server runs Oracle Database 19c and Oracle Linux 7.8 for the operating system.

This reference architecture document is for you if you are in one of the following roles:

- Database administrator
- Storage administrator
- Database performance analyzer
- IT professional with the responsibility of planning and deploying an Oracle Database solution

To use this reference architecture guide, you need familiarity with the following:

- Hitachi Virtual Storage Platform 5500
- Hitachi Advanced Server DS7000
- Hitachi Advanced Server DS120
- Storage area networks
- Oracle 19c Database
- Oracle Automatic Storage Management (Oracle ASM)
- Hitachi Adapters for Oracle Database
- Oracle Linux

Note – Testing of this configuration was in a lab environment. Many things affect production environments beyond prediction or duplication in a lab environment. Follow the recommended practice of conducting proof-of-concept testing for acceptable results in a non-production, isolated test environment that otherwise matches your production environment before your production implementation of this solution.

Solution Overview

This Hitachi Solution for Databases using Oracle Database is engineered, pre- tested, and qualified to provide high performance and high reliability in demanding, dynamic Oracle environments. It implements the Hitachi Unified Compute Platform with Oracle Database 19c leveraging Hitachi Virtual Storage Platform 5500 and the Hitachi Advanced Server DS7000. This reference architecture addresses the high performance and scalability requirements for OLTP and OLAP workloads.

Business Benefits

Here are some benefits of the new reference architecture:

- A high performance Oracle Database on Hitachi Virtual Storage Platform 5500
- Highly scalable Hitachi Advanced Server DS7000 series meet various performance and size requirements, while still leaving room for future growth
- Replace or consolidate existing IT infrastructure for Oracle databases with a smaller floor footprint, and smaller operating cost

High Level Infrastructure

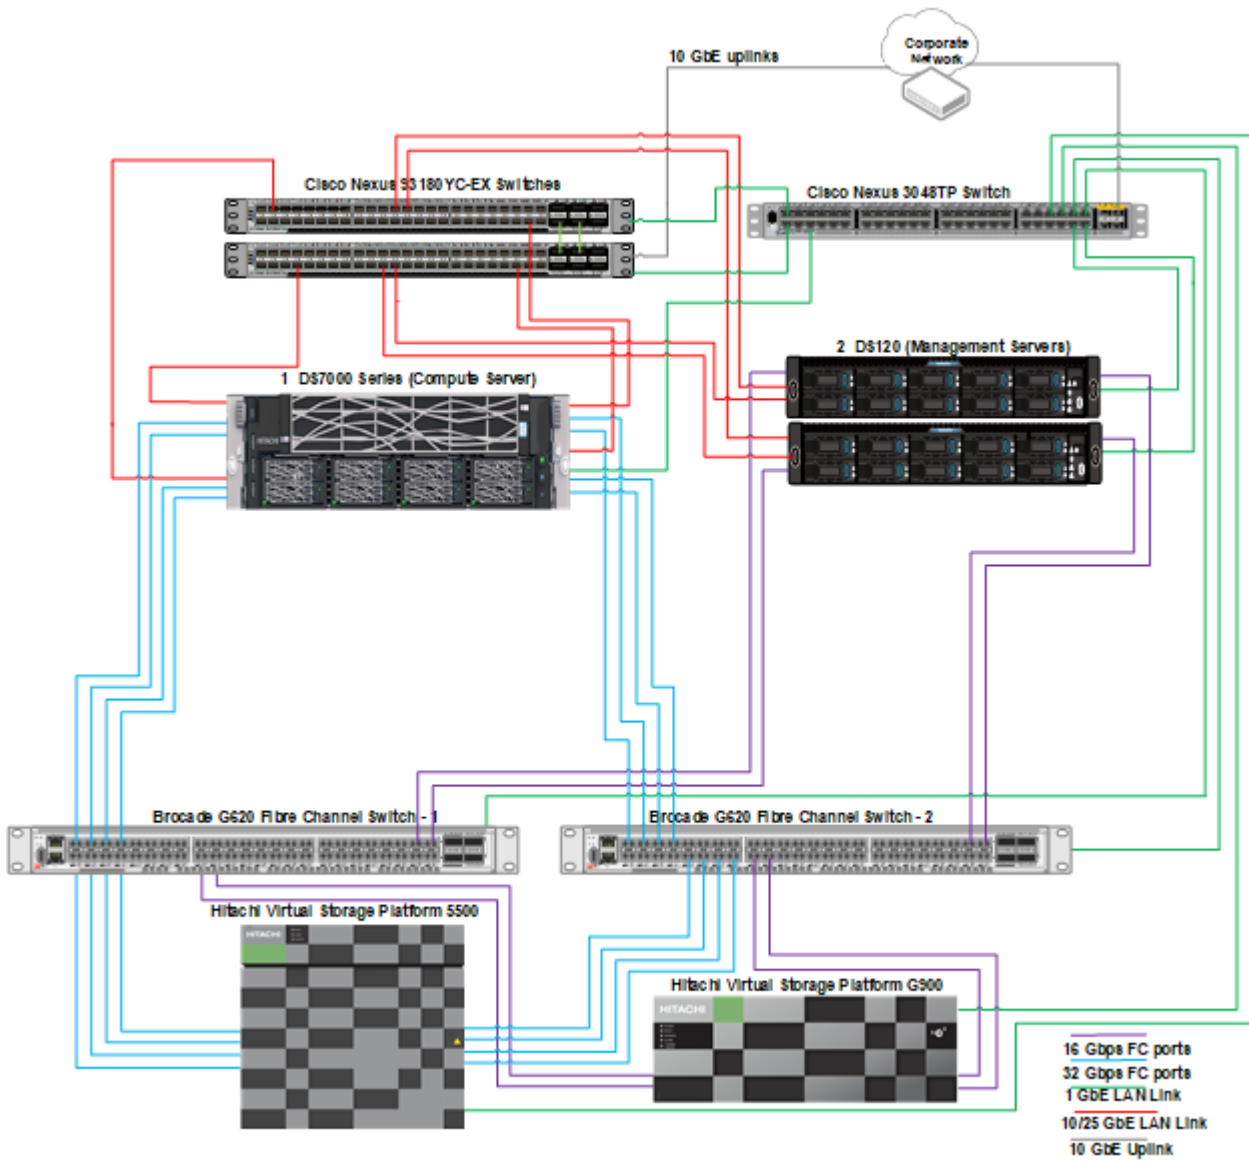
Hitachi Virtual Storage Platform 5500 and Hitachi Advanced Server DS7000 series are configured with fully redundant hardware for dual fabric connectivity between hosts and storage.

Figure 1 on page 3 is a high-level architecture diagram showing the following:

- Virtual Storage Platform 5500 and Hitachi Advanced Server DS7000 series server with a 4-socket CPU for a standalone Oracle 19c database configuration
- Two Hitachi Advanced Server DS120 units each with a 2-socket CPU for management of the architecture

Configuration of data volumes for the management servers is that used on Hitachi Virtual Storage Platform G900 during validation. Instead of Virtual Storage Platform G900, the management server in your implementation of this environment can be on any Virtual Storage Platform.

Figure 1



Hitachi Vantara recommends using a configuration with the following:

- A dedicated storage system for the production databases
- A dedicated storage system for data replication at a different site for business continuity, if needed

The uplink speed to the corporate network in Figure 1 was that used during validation. If higher bandwidth is required than 10/25 GbE, the Cisco Nexus 93180YC-EX switches can support uplink speeds of 40 GbE or 100 GbE.

Key Solution Components

The key solution components for this solution are listed in Table 1, "Hardware Components," on page 4, Table 2, "Software Components for Compute Nodes," on page 6, and Table 3, "Software Components for Management Nodes," on page 6.

TABLE 1. HARDWARE COMPONENTS

Hardware	Model	Detailed Description	Version	Quantity
Hitachi Virtual Storage Platform	5500	1 quad-controller block controller 4 CHA pairs (8 × 32 Gb/s Fibre Channel ports in use) 1 drive chassis with 4 drive trays 2048 GB cache memory 66 × 1.9 TB NVMe SSDs (two spare drives are recommended for production environments)	90-03-02-00/00	1
	G900	Two controllers 4 × 32 Gb/s Fibre Channel ports 16 × 12 Gb/s backend SAS ports 512 GB cache memory 4 × 6.0 TB 7.2K RPM SAS drives	88-02-03-60/00	1

TABLE 1. HARDWARE COMPONENTS (CONTINUED)

Hardware	Model	Detailed Description	Version	Quantity
Hitachi Advanced Server	DS7000 series	4 Intel Xeon Platinum 8276L CPUs @ 2.2 GHz	BIOS: BIOS_PUR043.37.02.104	1
		3072 GB (64 GB × 48) DIMM DDR4 synchronous registered (buffered), 2666 MHz	BMC: 33.45.00 (Build 445) Final Edition 20191114	
		4096 GB (256 GB × 16) Intel Optane Persistent Memory DIMMs*	M3PCPLD: 2.4.9.0 M3IOCPLD: 2.6.0.0	
	2 × Mellanox 10G MCX4121A-XCAT	Driver: mlnx-ofed Driver Version: 5.1-0.6.6.0 Firmware: 14.28.1300		
	4 Emulex LightPulse LPe32002-M2 2-Port 32 Gb/s Fibre Channel adapters	Driver: lpfc Driver Version: 12.0.0.10 Firmware: 12.0.261.9		
	2 Intel Xeon Silver 4110 CPUs @ 2.10 GHz	BIOS: S5BH3B16.H00	2	
256 GB (32 GB × 8) DIMM DDR4 synchronous registered (buffered) 2666 MHz	BMC: 4.64.06 CPLD:11			
1 × 64 GB MLC SATADOM for boot				
Brocade Fibre Channel Switches	G620	1 Dual port 25 GbE NIC Intel XXV710 PCIe card	Driver: i40e Driver Version: 1.7.17 Firmware: 6.8	2
		1 Emulex LightPulse LPe31002-M6 2-port 16 Gb/s Fibre Channel adapter	Firmware: 12.0.193.13 Driver: lpfc Driver version: 12.0.193.14	

TABLE 1. HARDWARE COMPONENTS (CONTINUED)

Hardware	Model	Detailed Description	Version	Quantity
Cisco Nexus switches	93180YC-EX	48 × 10/25 GbE fiber ports 6 × 40/100 Gb/s quad SFP (QSFP28) ports	BIOS: version 07.61 NXOS: version 7.0(3)I7(6)	2
	3048TP	1 GbE 48-Port Ethernet switch	BIOS: version 4.0.0 NXOS: version 7.0(3)I4(7)	1

Certain components may be optional, depending on the existing infrastructure and required interconnect topology. This may include the SAN, IP switches, and the management servers. However, this reference architecture documents the environment tested in the lab to support a full deployment of the architecture, including supporting components.

TABLE 2. SOFTWARE COMPONENTS FOR COMPUTE NODES

Software	Version	Function
Oracle Linux	7.8	Operating system
Oracle Database	19c	Database software
Oracle Grid Infrastructure	19c	Volume management, file system software, and Oracle Automatic Storage Management (ASM)
Oracle ASMLib	2.0.12	Oracle software for ASM

TABLE 3. SOFTWARE COMPONENTS FOR MANAGEMENT NODES

Software	Version	Function
VMware ESXi	6.7.0 Build 10302608	Hypervisor for management server
VMware vCenter Server	6.7.0 Build 10302608	VMware cluster management server
Hitachi Device Manager – Storage Navigator	Microcode dependent	Storage management software
Hitachi Ops Center	10.2.0	Hitachi infrastructure management software

Hitachi Virtual Storage Platform 5000 Series

This enterprise-class, flash array evolution storage, [Hitachi Virtual Storage Platform 5000](#) series (VSP) has an innovative, scale-out design optimized for NVMe and storage class memory. It achieves the following:

- **Agility using NVMe** – Speed, massive scaling with no performance slowdowns, intelligent tiering, and efficiency.
- **Resilience** – Superior application availability and flash resilience. Your data is always available, mitigating business risk.
- **Storage simplified** – Do more with less, integrate AI and ML (machine learning), simplify management, and save money and time with consolidation.

Hitachi Virtual Storage Platform 5500 is used in this reference architecture guide. Other models of the Hitachi Virtual Storage Platform 5000 series also support Oracle Real Application Clusters.

Hitachi Storage Virtualization Operating System RF

[Hitachi Storage Virtualization Operating System RF](#) (SVOS RF) spans and integrates multiple platforms. It integrates storage system software to provide system element management and advanced storage system functions. Used across multiple platforms, Storage Virtualization Operating System includes storage virtualization, thin provisioning, storage service level controls, dynamic provisioning, and performance instrumentation.

Hitachi Ops Center

Manage, optimize, orchestrate and protect your data with advanced IT analytics and automation using [Hitachi Ops Center](#). Achieve new insights, accelerate resource delivery, eliminate risks, and speed innovation to modernize your data center operations.

Use the power of AI operations with the following:

- [Analyzer](#). Improve IT operations with machine learning (ML) to drive resource service levels, utilization and automation at lower costs. Obtain operational visibility from virtual machines, servers, SAN switches to shared storage resources to optimize an application's full data path.
- [Automator](#). Deliver resources up to 70% faster than manual processes. Free staff to focus on strategic initiatives.
- [Protector](#). Meet tight service level requirements when protecting critical data and applications. Automatically support secondary business functions with data copies staff need to do their jobs. Make better use of backup data for activities, such as e-discovery and analysis. Simplify administration and replication management. Do it all with no disruption to production application availability and performance.

Hitachi Advanced Server DS7000 Series

To take advantage of the latest developments in artificial intelligence (AI), data analytics and machine learning, you require an infrastructure with high reliability, extreme performance, and agile scalability. [Hitachi Advanced Server 7000](#) series servers deliver this with a unique modular architecture.

Your server can be configured and scaled to meet the needs of a wide variety of application workloads. This can be used from in-memory data analytics processing to virtualization and hybrid cloud.

The Advanced Server DS7000 series has several complementary models, each based on the Intel Xeon scalable processor. You can upgrade a model to the next model, preserving your hardware and software investment as you grow.

Intel has launched the second-generation of Intel Xeon Scalable processors, revolutionary Intel Optane DC persistent memory, and new networking and storage options. Learn more about [Second Generation Intel Xeon Scalable Processors](#).

Hitachi Advanced Server DS120

Optimized for performance, high density, and power efficiency in a dual-processor server, [Hitachi Advanced Server DS120](#) delivers a balance of compute and storage capacity. This rack mounted server has the flexibility to power a wide range of solutions and applications.

The highly scalable memory supports up to 3 TB RAM using 24 slots of 2666 MHz DDR4 RDIMM. DS120 is powered by the Intel Xeon scalable processor family for complex and demanding workloads. There are flexible OCP and PCIe I/O expansion card options available.

The following applications were installed in individual virtual machines when testing this architecture. You would install them in most cases:

- VMware vCenter
- Hitachi Ops Center

Other management applications may be installed on additional virtual machines depending on your needs and requirements.

Oracle Linux

Using the stability and flexibility of [Oracle Linux](#), reallocate your resources to meet the next challenges instead of maintaining the status quo. Deliver meaningful business results by providing exceptional reliability of military-grade security. Use Oracle Linux to tailor your infrastructure as markets shift and technologies evolve.

Oracle Database

[Oracle Database](#) has a multi-tenant architecture so you can consolidate many databases quickly and manage them as a cloud service. Oracle Database also includes in-memory data processing capabilities for analytical performance. Additional database innovations deliver efficiency, performance, security, and availability. Oracle Database comes in two editions: Enterprise Edition and Standard Edition 2.

[Oracle Automatic Storage Management](#) (Oracle ASM) is a volume manager and a file system for Oracle database files. This supports single-instance Oracle Database and Oracle Real Application Clusters configurations. Oracle ASM is the recommended storage management solution that provides an alternative to conventional volume managers, file systems, and raw devices.

VMware vSphere

[VMware vSphere](#) is a virtualization platform that provides a datacenter infrastructure. It helps you get the best performance, availability, and efficiency from your infrastructure and applications. Virtualize applications with confidence using consistent management.

VMware vSphere has the following components:

- [VMware vSphere ESXi](#)

This hypervisor loads directly on a physical server. ESXi provides a robust, high-performance virtualization layer that abstracts server hardware resources and makes them shareable by multiple virtual machines.

- [VMware vCenter Server](#)

This provides a centralized platform for managing your VMware vSphere environments so you can automate and deliver a virtual infrastructure with confidence:

- VMware vSphere vMotion
- VMware vSphere Storage vMotion
- VMware vSphere Distributed Resource Scheduler
- VMware vSphere High Availability
- VMware vSphere Fault Tolerance

This reference architecture uses VMware vCenter Server for the management server cluster only.

Intel Optane Persistent Memory

[Intel Optane persistent memory](#) is an innovative memory technology that delivers a unique combination of affordable large capacity and support for data persistence. It can help businesses fuel innovation with increased capacity and unique memory modes, lower overall TCO while maximizing VM densities, and increase memory security with automatic hardware-level encryption.

We used memory mode for this solution.

Brocade Fibre Channel Switches from Broadcom

[Brocade Fibre Channel switches](#) deliver industry-leading performance, simplifying scale-out network architectures. Get the high performance, availability, and ease of management you need for a solid foundation to grow the storage network you want.

- [Broadcom G620](#) is a 1U-sized QSFP Fibre Channel switch with 48 × 16/32 Gb/s ports and 4 × 128 Gb/s ports for multiple-rack solutions.

Cisco Nexus Data Center Switches

[Cisco Nexus data center switches](#) are built for scale, industry-leading automation, programmability, and real-time visibility.

This solution uses the following Cisco switches to provide Ethernet connectivity:

- [Cisco Nexus 93180YC-EX](#) is a 1U-sized top-of-rack or leaf switch with 48 × 10/25 GbE ports (downlink) and 6 ports 40/100 GbE (uplink) for single-rack and multiple-rack solutions.
- [Cisco Nexus 3048TP](#) is a 1U-sized management switch with 48 ports 1 GbE for single-rack and multiple-rack solutions.

Solution Design

These are the reference architecture environments to implement Hitachi Unified Compute Platform for Oracle Database 19c. The environment uses one Hitachi Virtual Storage Platform 5500.

The infrastructure configuration includes the following:

- **Oracle Database Server** – One server is configured to run Oracle Database 19c.
- **Management Servers** – Two servers are used in the management server cluster.
- **Storage System** – There are virtual volumes mapped to each port that are presented to the server as LUNs.
- **SAN Connectivity** – There are SAN connections to connect the Fibre Channel HBA ports to the storage through Brocade G620 switches.

Although the configuration uses only one Oracle Database server, it is recommended to have at least two servers to provide the best availability and redundancy.

Storage Architecture

This is the storage architecture for this solution.

Storage Configuration

The design and deployment of the database storage configuration follows best practices from Hitachi Vantara and Oracle.

The high-level storage configuration diagram for this solution is shown in Figure 2 on page 11.

Figure 2

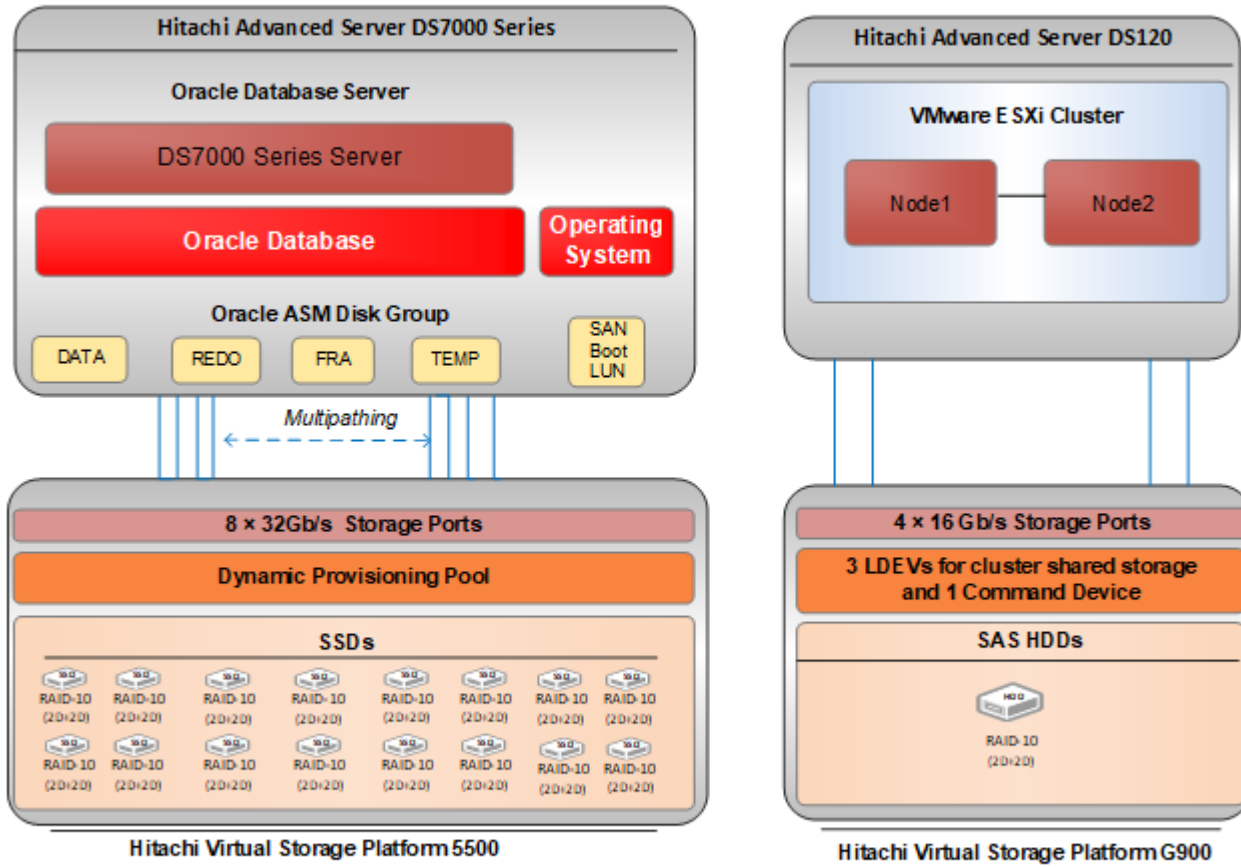


Table 4 shows the storage pool configuration used for this solution and lab verification.

TABLE 4. STORAGE POOL CONFIGURATION

Dynamic Provisioning Pool ID	Oracle-pool
RAID group	1-1 to 1-8, 2-1 to 2-8
RAID level	RAID-10 (2D+2D)
Drive type	1.9 TB SSD
Number of drives	64
Number of spare drives	2
Number of pool volume LDEVs	64
Pool volume LDEV size(s)	880.0 GB
Pool capacity	54.99 TB

Table 5 shows the Hitachi Virtual Storage Platform 5500 configuration for the operating systems and database volumes used in this solution.

TABLE 5. LOGICAL STORAGE CONFIGURATION

Dynamic Provisioning Pool ID	Oracle-pool					
Total number of DPVols	128	3	16	8	16	1
DPVol sizes (GB)	200	15	20	240	2000	380
ASM disk group	DATA	OCR	REDO	TEMP	FRA	N/A
Purpose	OLTP application tablespaces System Sysaux Undo	Oracle cluster Registry (GI installation)	Online redo Logs Control files	Temp	Incremental backups Archived redo Logs Control file Auto backups	SAN boot OS volumes
Storage port	1A, 1B, 1C, 1D, 2A, 2B, 2C, 2D					

Table 6 shows the Hitachi Virtual Storage Platform G900 configuration for management servers.

TABLE 6. HITACHI VIRTUAL STORAGE PLATFORM G900 CONFIGURATION FOR MANAGEMENT SERVERS

Item	Value or Description
Purpose	VMware Datastores Command control interface device
RAID level	RAID-10 (2D+2D)
Drive type	6.0 TB 7.2K RPM SAS
Number of drives	4
Number of spare drives	0
Number of LDEVs	3
LDEV size(s)	3000 GB
Number and size of CCI devices	1 × 100 MB
Storage port for management servers	7A, 7B, 8A, 8B

An additional RAID group consisting of four 3 TB, 7.2k RPM SAS drives configured as RAID-10 (2D+2D) was used as shared storage for the management server cluster. A 3 TB LUN and a command device were mapped to four storage ports.

Additional LUNs can be mapped if required. While the test environment was configured using a dedicated SAS RAID group for the management server cluster, this can be configured as a dedicated SSD RAID group, a dedicated dynamic provisioning pool, or it can use capacity on the dynamic provisioning pool configured for the Oracle environment, depending on your requirements.

Database Layout

The database layout design follows recommended best practices from Hitachi Vantara for Hitachi Virtual Storage Platform 5500 for small random I/O traffic, such as OLTP transactions. The layout also follows Oracle ASM best practices when using the Hitachi storage system and Hitachi Advanced Server.

The design can vary greatly from one implementation to another, based on the RAID configuration and number of drives used during the implementation. The components in this solution set have the flexibility for use in various deployment scenarios to provide the right balance between performance and ease of management for a given scenario. HBAs are installed to both odd and even number PCIe slots on each Advanced Server DS7000 module to provide balanced I/O paths to the storage system.

Oracle Configuration

Configure Oracle for this solution doing the following:

- **Data and Indexes Tablespace** – Assign an ASM diskgroup with external redundancy for the data and index tablespaces.
- **TEMP Tablespace** – Place the TEMP tablespace in the TEMP ASM diskgroup.
- **Undo Tablespace** – Create an UNDO tablespace within the Oracle Data ASM diskgroup.
- **Online Redo Logs** – Create an ASM diskgroup with external redundancy for Oracle online redo logs.
- **Oracle Cluster Registry** – Create a disk group for the Oracle Database 19c non-cluster GI installation.

Oracle Parameters

Table 7 shows Oracle ASM and Oracle Database parameters.

TABLE 7. ORACLE ASM AND DATABASE PARAMETERS

Category	Item	Value
Oracle RAC option	RAC configuration	No
	ASM	Yes
Oracle ASM environment parameters	OCR	3 × 15 GB
	DATA	128 × 200 GB
	REDO	16 × 20 GB
	TEMP	8 × 240 GB
	FRA	16 × 2,000 GB
Oracle Database environment parameters	SGA_TARGET	1,536 GB
	PGA_AGGREGATE_TARGET	768 GB
	DB_CACHE_SIZE	768 GB
	DB_KEEP_CACHE_SIZE	192 GB
	DB_RECYCLE_CACHE_SIZE	192 GB
	LOG_BUFFER	512 MB
	USE_LARGE_PAGES	TRUE
	FILESYSTEMIO_OPTIONS	SETALL
	DISK_ASYNCH_IO	TRUE

Oracle ASM Disk Mappings

Table 8 shows volumes, LUNs, and ASM disk groups.

TABLE 8. VOLUMES, LUNS, AND ASM DISKGROUPS

ASM Disk Group	ASM Disk	DM-Multipath LUNs	LUN Details	Purpose
OCR	OCR1 - OCR3	/dev/mapper/mpathaa - /dev/mapper/mpathac	3 × 15 GB	Oracle 19c GI installation
REDO	REDO01 - REDO16	/dev/mapper/mpathca - /dev/mapper/mpathcc /dev/mapper/mpathbn - /dev/mapper/mpathbz	16 × 20 GB	Online REDO log group
FRA	FRA01 - FRA16	/dev/mapper/mpathcl - /dev/mapper/mpathco	16 × 2000 GB	Flash recovery area
TEMP	TEMP1 - TEMP8	/dev/mapper/mpathcd - /dev/mapper/mpathck	8 × 240 GB	Temp
DATA	DATA001 - DATA128	/dev/mapper/mpathdd - /dev/mapper/ mpathdp /dev/mapper/mpatheb - /dev/mapper/mpathep /dev/mapper/mpathfa - /dev/mapper/mpathfp /dev/mapper/mpathga - /dev/mapper/mpathgp	128 × 200 GB	Application data

Oracle Server Configuration

Table 9 shows the operating system configurations for the Oracle servers.

TABLE 9. ORACLE SERVER OPERATING SYSTEM CONFIGURATIONS

Server Configuration	Server Operating System Setting Details
RPMs for Oracle Database 19c	binutils-2, compat-libcap1-1, compat-libstdc++-33-3, gcc-4, gcc-c++-4, glibc-2, glibc-devel-2, ksh, libaio-0, mlbaio-devel-0, libgcc-4, libstdc++-4, libstdc++-devel-4, libXi-1, libXtst-1, make-3, sysstat-10
/etc/multipath.conf	user_friendly_names yes find_multipaths yes path_grouping_policy multibus path_selector "service-time 0"
/etc/security/limits.conf	oracle soft nfile 1024 oracle hard nfile 65536 oracle soft nproc 16384 oracle hard nproc 16384 oracle soft stack 10240 oracle hard stack 32768 oracle hard memlock 356537484 oracle soft memlock 356537484 grid hard nfile 65536

TABLE 9. ORACLE SERVER OPERATING SYSTEM CONFIGURATIONS (CONTINUED)

Server Configuration	Server Operating System Setting Details
/etc/sysctl.conf	fs.aio-max-nr = 6291456 fs.file-max = 6815744 kernel.shmall = 2097152 kernel.shmmax = 536870912 kernel.shmmni = 4096 kernel.sem = 250 32000 100 128 net.ipv4.ip_local_port_range = 9000 65500 net.core.rmem_default = 262144 net.core.rmem_max = 4194304 net.core.wmem_default = 262144 net.core.wmem_max = 1048586 vm.nr_hugepages = 1000000
Swap space	64 GB

For an OLTP database, Hitachi Vantara recommends setting "service-time 0" for path_selector in /etc/multipath.conf file for the best performance.

Management Server Configuration

Table 10 shows the Hitachi Virtual Storage Platform G900 configuration for the management server cluster.

TABLE 10. HITACHI VSP G900 MANAGEMENT SERVER CLUSTER CONFIGURATION

Attribute	Value
Purpose	VMware shared datastores Command device
RAID level	RAID-10 (2D+2D)
Drive type	3 TB 7.2K RPM SAS
Number of drives	4
Number of spare drives	0
Number and size of LDEVs for datastores	1 × 3000 GB

TABLE 10. HITACHI VSP G900 MANAGEMENT SERVER CLUSTER CONFIGURATION (CONTINUED)

Attribute	Value
Number and size of command devices	1 × 100 MB
Storage port for management servers	7A, 7B, 8A, 8B

Server and Application Architecture

This reference architecture uses four Hitachi Advanced Server DS7000 series servers with second-generation Intel Xeon scalable processors for each storage system architecture that was tested.

This provides the compute power for the Oracle database to handle complex database queries and a large volume of transaction processing in parallel.

Two Hitachi Advanced Server DS120 servers are used for VMware ESXi management server configuration.

Table 11 describes the details of the server configurations for this solution.

TABLE 11. HITACHI ADVANCED SERVER DS7000 SERIES SERVER AND ADVANCED SERVER DS120 SPECIFICATIONS

Hitachi Advanced Server	Server	Host Name	Role	CPU Core	RAM
DS7000	Oracle Server	orcl	Oracle database server	112 (4 × 28C)	3,072 GB (64 GB × 48) or 1,536 GB (64 GB × 24) + 4,096 GB ((256 GB × 16) Intel Optane Persistent Memory)
DS120	Management Server 1	VMware ESXi 1	Management virtual machine hosts	16 (2 × 8C)	256 GB (32 GB × 8)
	Management Server 2	VMware ESXi 2		16 (2 × 8C)	256 GB (32 GB × 8)

SAN Architecture

Map the provisioned LDEVs to multiple ports on each Hitachi Virtual Storage Platform 5500. These LDEV port assignments provide multiple paths to the storage system from the host for high availability. This reference architecture uses four dual port Emulex HBAs on the Hitachi Advanced Server DS7000 series server.

- 8 SAN switch connections for Virtual Storage Platform 5500 Fibre Channel ports
- 8 SAN switch connections for Advanced Server DS7000 HBA ports
- 4 SAN switch connections for Virtual Storage Platform G900 Fibre Channel ports
- 4 SAN switch connections for management server HBA ports

Table 12 shows details of the Fibre Channel switch connect configuration on Hitachi Virtual Storage Platform 5500, Virtual Storage Platform G900, and database servers.

TABLE 12. HITACHI VIRTUAL STORAGE PLATFORM 5500, VIRTUAL STORAGE PLATFORM G900, SERVERS, AND FIBRE CHANNEL SWITCH CONFIGURATIONS

Server	HBA Ports	Storage Host Group	Switch Zone	Storage System	Storage Port	Brocade G620 Switch
DS7000	HBA1_1	DS7000_59_HBA1_1	DS7000_59_HBA1_1_ASE47_108_1A	VSP 5500	1A	SAN-switch 1
	HBA1_2	DS7000_59_HBA1_2	DS7000_59_HBA1_2_ASE47_108_2A		2A	SAN-switch 2
	HBA2_1	DS7000_59_HBA2_1	DS7000_59_HBA2_1_ASE47_108_1B		1B	SAN-switch 1
	HBA2_2	DS7000_59_HBA2_2	DS7000_59_HBA2_2_ASE47_108_2B		2B	SAN-switch 2
	HBA3_1	DS7000_59_HBA3_1	DS7000_59_HBA3_1_ASE47_108_1C		1C	SAN-switch 1
	HBA3_2	DS7000_59_HBA3_2	DS7000_59_HBA3_2_ASE47_108_2C		2C	SAN-switch 2
	HBA4_1	DS7000_59_HBA4_1	DS7000_59_HBA4_1_ASE47_108_1D		1D	SAN-switch 1
	HBA4_2	DS7000_59_HBA4_2	DS7000_59_HBA4_2_ASE47_108_2D		2D	SAN-switch 2
DS120 Server 1	HBA1_1	DS120_68_HBA1_1	DS120_68_HBA1_1_ASE32_50_7A	VSP G900	7A	SAN-switch 1
	HBA1_2	DS120_68_HBA1_2	DS120_68_HBA1_2_ASE32_50_8A		8A	SAN-switch 2
DS120 Server 2	HBA1_1	DS120_69_HBA1	DS120_69_HBA1_1_ASE32_50_7B		7B	SAN-switch 1
	HBA1_2	DS120_69_HBA2	DS120_69_HBA1_2_ASE32_50_8B		8B	SAN-switch 2

Note – In a production environment without a storage system, such as Hitachi Virtual Storage Platform G900 for management servers, separate storage ports on the Virtual Storage Platform 5000 series are recommended for the management servers to avoid impact on the database performance. While shared storage ports can be used, monitor port utilization to avoid performance issues in extreme high performance environments.

Network Architecture

This architecture requires the following separate networks:

- **Public Oracle Network** – This network provides client connections to the Oracle database and other applications.
- **Management Network** – This network is for hardware management console connections.

Hitachi Vantara recommends using pairs of 25 Gb/s NICs for the public network.

Observe these points when configuring public networks in your environment:

- Use NIC bonding to provide failover and load balancing of interconnections within a server.
- Set all NICs to full duplex mode.
- Use at least two public NICs for client connections to the application and database.

Physical Network Configuration

Figure 3 shows the IP network switch connection.

Figure 3

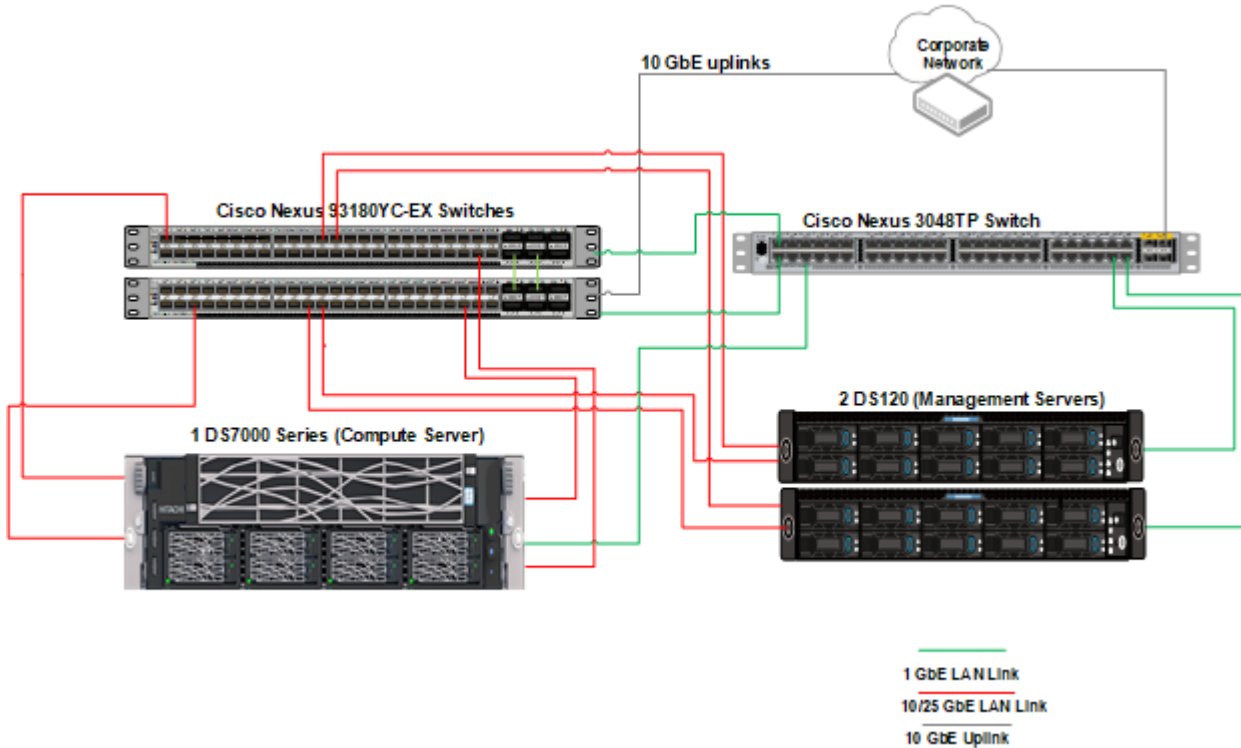


Table 13 shows the network configuration of the database server and Table 14, "Management Server IP Network Configuration," on page 21 shows the network configuration of management servers used when testing the environment. Your values may be different.

Configure NIC bonding or NIC teaming with NIC ports across multiple physical NICs to avoid a single point of failure.

TABLE 13. DATABASE SERVER NETWORK CONFIGURATION

Server	NIC Port	VLAN or Subnet	NIC Bond	IP Address	Network	Bandwidth (Gb/s)	Cisco Nexus 93180YC-EX Switch	
							Switch Number	Port
Oracle Database Server	NIC 1 PORT 1	242	Bond0	10.76.41.169	Public Oracle	10	1	31
	NIC 2 PORT 1					10	2	32
	BMC- Dedicated NIC	242	-	10.76.41.167	Management	1	-	-

TABLE 14. MANAGEMENT SERVER IP NETWORK CONFIGURATION

Server	NIC Port	VLAN or Subnet	NIC Bond	IP Address	Network	Bandwidth (Gb/s)	Cisco Nexus 93180YC-EX Switch	
							Switch Number	Port
DS120 Server 1	NIC - 0	242	ESXi NIC Teaming	10.76.32.80	Public	25	1	39
	NIC - 1					25	2	40
	BMC- Dedicated NIC	242	-	10.76.32.169	Server Management	1	-	
DS120 Server 2	NIC - 0	242	ESXi NIC Teaming	10.76.32.81	Public	25	1	41
	NIC - 1					25	2	42
	BMC- Dedicated NIC	242	-	10.76.32.170	Server Management	1	-	

Table 15 lists the virtual machine configuration running on the management server cluster. The virtual machine configurations in Table 15 are required for the versions used for testing in the lab.

TABLE 15. MANAGEMENT SERVER CLUSTER VIRTUAL MACHINE CONFIGURATION

Virtual Machine	vCPU	Virtual Memory	Disk Capacity	IP Address
Hitachi Oracle Adapters	2	6 GB	40-50 GB	192.168.242.80
Hitachi Ops Center Administrator	4	16 GB	100 GB	192.168.242.81
Hitachi Ops Center Analyzer	4	32 GB	800 GB	192.168.242.194
Hitachi Ops Center Analyzer detail view	4	10 GB	110 GB	192.168.242.197
vCenter	2	10 GB	300 GB	192.168.242.102

Engineering Validation

This summarizes the key observations from the test results for the Hitachi Unified Compute Platform architecture to deploy Oracle Database 19c with Hitachi Virtual Storage Platform 5500.

Test Methodology

This describes the methodology used for testing.

Oracle Orion

Oracle Orion is a tool for predicting the performance of an Oracle database without having to install Oracle or create a database. Unlike other I/O calibration tools, Oracle Orion is expressly designed for simulating Oracle database I/O workloads using the same I/O software stack as Oracle. Orion can also simulate the effect of striping performed by Oracle Automatic Storage Management.

For more information about Orion, see "I/O Configuration and Design" in the [Oracle Database Performance Tuning Guide](#).

peakmarks

The test results are created using peakmarks OLTP test cases.

For Oracle platforms that is used for the following purposes, peakmarks is benchmark software:

- Performance verification (quality assurance)
- Evaluation of different infrastructure products, technologies, and solutions (price/performance comparison)
- Performance optimization (improvement in efficiency)

This provides transparency and comparability in price versus performance considerations for Oracle infrastructures. The peakmarks 9.4 build 20191130 tool was used to validate this solution.

Database Configuration

Table 16 shows parameter details for the standalone Oracle Database 19c.

TABLE 16. STANDALONE ORACLE DATABASE 19C PARAMETERS

Oracle Database Parameter	Value
compatible	19.3.0.0.0
cluster_database	FALSE
cluster_database_instances	1
Oracle Database size	16 TB
Database storage type	ASM
Database fill factor	80%

Test Environment

Table 17 lists the test environment.

TABLE 17. TEST ENVIRONMENT DETAILS

Item	Description	Value	
		Without Intel Optane Persistent Memory	With Intel Optane Persistent Memory
Operating system		OL 7.8 with UEK 4.14.35-1902.300.11.el7uek.x86_64	OL 8.2 with UEK 5.4.17_2011.6.2.el8uek.x86_64
Workload type		OLTP/OLAP	
Database size		16 TB	
Number of physical cores	CPU cores	112	
DRAM Memory	Server memory	3,072 GB	1,536 GB
Intel Optane Persistent Memory	Server memory	0	4,096 GB
Network	Server IP network	2 × 25 Gb/s NIC bonding	

Test Results

Table 18 shows results of the Orion and peakmarks tests.

TABLE 18. ORION AND PEAKMARKS TEST RESULTS

Test Category	Test Case	Test Results			
		With 48 × 64 GB DRAM DIMMs		With 24 × 64 GB DRAM DIMMs and 16 × 256 GB AEP DIMMs	
		Throughput	RT (ms)	Throughput	RT (ms)
ORION	8KB Random Read	1,561,006 IOPS	0.83	N/A	N/A
	8KB Random Write	604,840 IOPS	0.81	N/A	N/A
	1MB Sequential Read	26.84 GB/s	N/A	N/A	N/A
	1MB Sequential Write	16.65 GB/s	N/A	N/A	N/A

TABLE 18. ORION AND PEAKMARKS TEST RESULTS (CONTINUED)

Test Category	Test Case	Test Results			
		With 48 × 64 GB DRAM DIMMs		With 24 × 64 GB DRAM DIMMs and 16 × 256 GB AEP DIMMs	
		Throughput	RT (ms)	Throughput	RT (ms)
peakmarks	STO-READ	24.91 GB/s	3.90	23.79 GB/s	7.28
	STO-Random 100% read	1,238,452 IOPS	0.96	1,252,623	0.95
	STO-Random 80% read/20% write	1,363,887 IOPS	0.87	1,539,557	0.98
	STO-Random 70% read/30% write	1,561,952 IOPS	1.11	1,701,044	0.92
	STO-Random 100% write	891,794 IOPS	0.38	984,657	0.73
	SRV-SCAN	113.04 GB/s	N/A	133.51 GB/s	N/A
	DBWR-THR	5.55 GB/s	N/A	4.98 GB/s	N/A
	DA-LOW/ROW	109.69 GB/s	N/A	134.80 GB/s	N/A
	DA-LOW/STO	30.89 GB/s	N/A	31.02 GB/s	N/A
	DA-LOW/COL	1,645.95 GB/s	N/A	1,673.64 GB/s	N/A
	TP-LIGHT 1 RPT 0% update	420,548 tps	0.26	308,917 tps	0.41
	TP-LIGHT 1 RPT 20% update	277,695 tps	1.00	268,607 tps	0.62
	TP-LIGHT 1 RPT 30% update	280,288 tps	0.99	295,659 tps	0.75
	TP-MEDIUM 25 RPT 0% update	66,874 tps	4.15	21,567 tps	5.13
	TP-MEDIUM 25 RPT 20% update	14,729 tps	11.35	12,159 tps	8.36
TP-MEDIUM 25 RPT 30% update	13,818 tps	8.07	11,818 tps	8.19	

Conclusion

Hitachi UCP for Oracle databases solutions are well designed and balanced for different workloads with minimal or no tuning effort needed when workloads shift from one to another. Our test results clearly show that a single module Hitachi VSP 5500 with SSDs can deliver superior Oracle database performance with OLTP and OLAP test cases at the level we have seen with Hitachi FMDs. Intel Optane Persistent Memory provides a larger memory configuration option to the database servers with lower cost. We believe the best use case of the Intel Optane Persistent Memory is for OLAP workloads until Oracle database software supports App Direct Mode or Mixed Memory Mode.

For More Information

Hitachi Vantara Global Services offers experienced storage consultants, proven methodologies and a comprehensive services portfolio to assist you in implementing Hitachi products and solutions in your environment. For more information, see the [Services](#) website.

Demonstrations and other resources are available for many Hitachi products. To schedule a live demonstration, contact a sales representative or partner. To view on-line informational resources, see the [Resources](#) website.

Hitachi Academy is your education destination to acquire valuable knowledge and skills on Hitachi products and solutions. Our Hitachi Certified Professional program establishes your credibility and increases your value in the IT marketplace. For more information, see the Hitachi Vantara [Training and Certification](#) website.

For more information about Hitachi products and services, contact your sales representative, partner, or visit the [Hitachi Vantara](#) website.

Hitachi Vantara



Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA
www.HitachiVantara.com | community.HitachiVantara.com

Regional Contact Information
USA: 1-800-446-0744
Global: 1-858-547-4526
HitachiVantara.com/contact

© Hitachi Vantara LLC 2020. All rights reserved. HITACHI is a trademark or registered trademark of Hitachi, Ltd. VSP is a trademark or registered trademark of Hitachi Vantara LLC. All other trademarks, service marks and company names are properties of their respective owners.

Notice: This document is for informational purposes only, and does not set forth any warranty, expressed or implied, concerning any equipment or service offered or to be offered by Hitachi Vantara.

MK- SL-203-01. October 2020.