

Hitachi Solution for Databases - Oracle Real Application Clusters Database 19c with Hitachi VSP 5600 and Hitachi Advanced Server DS220 G2

Reference Architecture Guide

© 2022 Hitachi Vantara LLC. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including copying and recording, or stored in a database or retrieval system for commercial purposes without the express written permission of Hitachi, Ltd., or Hitachi Vantara LLC (collectively "Hitachi"). Licensee may make copies of the Materials provided that any such copy is: (i) created as an essential step in utilization of the Software as licensed and is used in no other manner; or (ii) used for archival purposes. Licensee may not make any other copies of the Materials. "Materials" mean text, data, photographs, graphics, audio, video and documents.

Hitachi reserves the right to make changes to this Material at any time without notice and assumes no responsibility for its use. The Materials contain the most current information available at the time of publication.

Some of the features described in the Materials might not be currently available. Refer to the most recent product announcement for information about feature and product availability, or contact Hitachi Vantara LLC at https://support.hitachivantara.com/en_us/contact-us.html.

Notice: Hitachi products and services can be ordered only under the terms and conditions of the applicable Hitachi agreements. The use of Hitachi products is governed by the terms of your agreements with Hitachi Vantara LLC.

By using this software, you agree that you are responsible for:

1. Acquiring the relevant consents as may be required under local privacy laws or otherwise from authorized employees and other individuals; and
2. Verifying that your data continues to be held, retrieved, deleted, or otherwise processed in accordance with relevant laws.

Notice on Export Controls. The technical data and technology inherent in this Document may be subject to U.S. export control laws, including the U.S. Export Administration Act and its associated regulations, and may be subject to export or import regulations in other countries. Reader agrees to comply strictly with all such regulations and acknowledges that Reader has the responsibility to obtain licenses to export, re-export, or import the Document and any Compliant Products.

Hitachi and Lumada are trademarks or registered trademarks of Hitachi, Ltd., in the United States and other countries.

AIX, AS/400e, DB2, Domino, DS6000, DS8000, Enterprise Storage Server, eServer, FICON, FlashCopy, GDPS, HyperSwap, IBM, Lotus, MVS, OS/390, PowerHA, PowerPC, RS/6000, S/390, System z9, System z10, Tivoli, z/OS, z9, z10, z13, z14, z/VM, and z/VSE are registered trademarks or trademarks of International Business Machines Corporation.

Active Directory, ActiveX, Bing, Excel, Hyper-V, Internet Explorer, the Internet Explorer logo, Microsoft, Microsoft Edge, the Microsoft corporate logo, the Microsoft Edge logo, MS-DOS, Outlook, PowerPoint, SharePoint, Silverlight, SmartScreen, SQL Server, Visual Basic, Visual C++, Visual Studio, Windows, the Windows logo, Windows Azure, Windows PowerShell, Windows Server, the Windows start button, and Windows Vista are registered trademarks or trademarks of Microsoft Corporation. Microsoft product screen shots are reprinted with permission from Microsoft Corporation.

All other trademarks, service marks, and company names in this document or website are properties of their respective owners.

Copyright and license information for third-party and open source software used in Hitachi Vantara products can be found in the product documentation, at <https://www.hitachivantara.com/en-us/company/legal.html> or https://knowledge.hitachivantara.com/Documents/Open_Source_Software.

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

Changes	Date
Updated peakmarks information.	November 8, 2022
Initial release	February 28, 2022

Reference Architecture Guide

Use this reference architecture guide to understand how Hitachi Solution for Databases provides a high-performance, low latency, integrated, converged solution for Oracle using Hitachi Virtual Storage Platform 5600 (VSP 5600). The environment uses VSP 5600, Hitachi Advanced Server DS220 G2, and DS120 G2 servers with Ice Lake Server Scalable Processors. With these products, design an Oracle converged infrastructure to meet your requirements and budget.

This solution uses the high-performing Hitachi Advanced Server DS220 G2 to boost performance and lower I/O latency. The VSP 5600 storage system with NVMe SSDs is used in this reference architecture to run a dedicated Oracle Database 19c with the Oracle Real Application Clusters (RAC) option, and it uses Oracle Linux 8.5 for the operating system. Hitachi Advanced Server DS120 G2 is dedicated for management applications. This 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 must be familiar with the following:

- Hitachi Virtual Storage Platform 5600
- Hitachi Advanced Server DS220 G2
- Hitachi Advanced Server DS120 G2
- Storage area networks
- Oracle 19c RAC Database
- Oracle Automatic Storage Management (Oracle ASM)
- 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

The Hitachi solution for Oracle RAC databases on Hitachi Unified Compute Platform (UCP) for Oracle Database is engineered, pre-tested, and qualified to provide high performance and high reliability in demanding and dynamic Oracle environments.

This reference architecture implements Hitachi Unified Compute Platform for Oracle Real Application Clusters on two nodes using Hitachi Virtual Storage Platform 5600. It addresses the high availability, performance, and scalability requirements for OLTP and OLAP workloads. This solution was developed using Intel Whitley Platform-based Hitachi Advanced Server DS220 G2 servers, DS120 G2 servers, and VSP E990 storage for the management server environment.

Business benefits

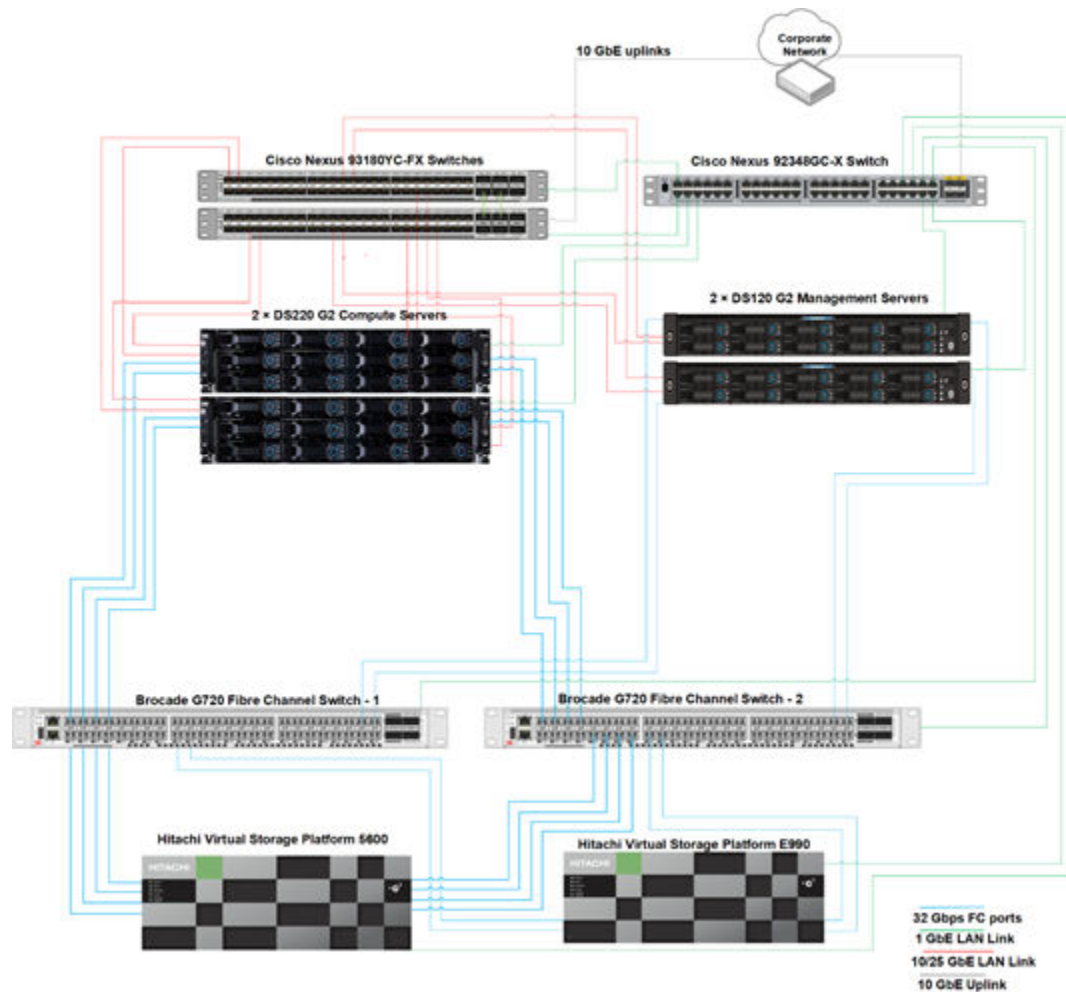
Here are some benefits of this reference architecture:

- Achieves high Oracle Database performance with VSP 5600.
- Provides a solution for customers who are looking for very low I/O latency for Oracle RAC database.

High level infrastructure

VSP 5600 and Hitachi Advanced Server DS220 G2 were configured with fully redundant hardware for dual fabric connectivity between hosts and storage.

The following figure shows the high-level architecture diagram using VSP 5600 and Hitachi Advanced Server DS220 G2 for a two-node Oracle 19c RAC configuration, and two Hitachi Advanced Server DS120 G2 servers for management of the architecture. Data volumes for the management servers are configured on a VSP E990, which is shared across different management nodes. The VSP E990 is not required in your environment; data volumes can also be configured on the VSP 5600 or other VSP storage systems.



To avoid a performance impact to the production database, Hitachi Vantara recommends using a configuration with the following:

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

Hitachi Vantara recommends using pairs of 25 Gbps NICs for the cluster interconnect network and public network.

Uplink speed to the corporate network depends on the customer environment and requirements. The Cisco Nexus 93180YC-FX switches can support uplink speeds of 40 GbE or 100 GbE if higher bandwidth is required.

Key solution components

The key components for this solution are listed in the following tables.

Table 1 Hardware Components

Vendor	Hardware	Detailed Description	Version	Quantity
Hitachi Vantara	Hitachi Virtual Storage Platform 5600	4 × CHB pairs 8 × 32 Gbps Fibre Channel ports 2048 GB cache memory 64 × 3.84 TB backend NVMe SSDs	90-08-01-00/00	1
	Hitachi Virtual Storage Platform E990	Two Controllers 4 × 32 Gbps Fibre Channel ports 16 × 12 Gbps backend SAS ports 1024 GB cache memory 17 × 3.8 TB NVMe SSDs*	93_02_02_60/00	1
Hitachi Vantara	Hitachi Advanced Server DS220 G2	2 × Intel Xeon Platinum 8368 38C CPU @ 2.40GHz 768 GB (64 GB × 12) DIMM DDR4-3200 Synchronous Registered (Buffered) 3200 MHz	BIOS: S5XH3A12.H03 BMC: 3.16.06 CPLD: 07	2
		2 Intel E810 dual port 25 GbE NIC cards	Driver: ice Driver Version: 0.8.2-k Firmware: 2.42	
		2 × Emulex LightPulse LPe35002-M2 2-Port 32 Gb Fibre Channel Adapters	Driver: lpfc Driver Version: 12.8.0.10 Firmware: 12.8.542.26	
Hitachi Advanced Server DS120 G2	2 × Intel Xeon Processor 4310, 12-core, 2.10 GHz, 120W	BIOS: S5XH3A12.H03 BMC: 3.16.06	2	

Vendor	Hardware	Detailed Description	Version	Quantity
		256 GB (32 GB × 8) DIMM DDR4-3200 Synchronous Registered (Buffered) 3200 MHz 1 × 256 GB NVMe 0.3DWPD M.2 SSD for boot	CPLD: 07	
		1 × Dual Port 25 GbE NIC Intel E810 PCIe card	Driver Version: ice 0.8.2-k Firmware: 2.42	
		1 × Emulex LightPulse LPe35002-M2 2-Port 32Gb Fibre Channel Adapters	Driver: lpfc Driver Version: 12.8.0.10 Firmware: 12.8.542.26	
Brocade	G720 Fibre Channel switches	48 × 32Gbps ports Fibre Channel switch 32 Gbps SFPs	Fabric OS : v9.0.1c	2
Cisco	Nexus 93180YC-FX	48 × 10/25 GbE ports 6 × 40/100 Gbps Quad SFP (QSFP28) ports	NXOS: version 9.3.7	2
	Cisco Nexus 92348TP	48 × 1 GbE ports Ethernet Switch	NXOS: version 9.3.7	1

*Customers can choose larger capacity SSDs to fit their business requirements.

Certain components may be optional depending on the existing infrastructure and required interconnect topology. This might 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

Software	Version	Function
Oracle Linux	8.5 with UEK kernel 6.3	Operating system
Oracle Database	19.12.0.0.0	Database software

Software	Version	Function
Oracle Real Application Cluster	19.12.0.0.0	Cluster software
Oracle Grid Infrastructure	19.12.0.0.0	Volume management, file system software, and Oracle Automatic Storage Management (ASM)
Oracle ASMLib	2.1.12	Oracle software for ASM

Table 3 Management Node Software Components

Software	Version	Function
ESXi	7.0 U2	Hypervisor for management server
vCenter Server	7.0 U2	VMware cluster management server
Hitachi Device Manager - Storage Navigator	Microcode dependent	Storage management software
Hitachi Ops Center	10.8.0	Hitachi infrastructure management software

Solution design

This section describes the reference architecture environments used to implement Hitachi Unified Compute Platform (UCP) for Oracle with the Real Application Clusters option. Each environment uses one Hitachi Virtual Storage Platform 5600.

The infrastructure configuration includes the following:

- Oracle RAC Servers – There are two server nodes configured in an Oracle Real Application Cluster.
- Management Servers – Two servers are used in the management server cluster.
- Storage System – There are vVols mapped to each port that are presented to the server as LUNs.
- SAN Connection – There are SAN connections to connect the Fibre Channel HBA ports to the storage through Brocade G720 switches.

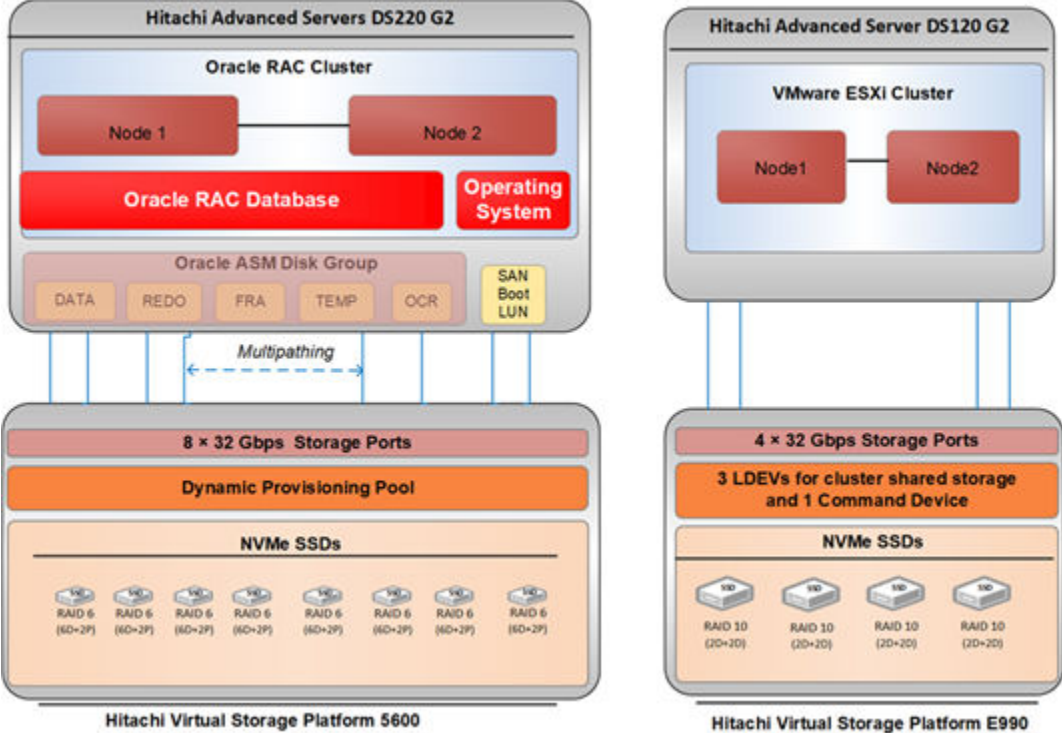
Storage architecture

This section describes the storage architecture for this solution.

Storage configuration

The storage configuration takes into consideration Hitachi Vantara for Hitachi Virtual Storage Platform and Oracle recommended best practices for the design and deployment of database storage.

The high-level storage configuration diagram for this solution is shown in the following figure.



The following table lists the storage pool configuration used for this solution and lab verification.

Table 4 Storage Pool Configuration

Pool ID	Oracle(0)
Pool Type	Dynamic Provisioning Pool
RAID Group	1-1 to 1-8
RAID Protection Type	RAID 6 (6D+2P)
Drive Type/Capacity	3.84 TB NVMe SSD
Number of Drives	64
Number of LDEVs	96
LDEV Size(s)	1760 GB
Pool Capacity	164.99 TB

The following table lists the logical storage configuration used in this solution.

Table 5 Logical Storage Configuration

Item	Value/Description					
Total number of DPVols	128	3	32	16	16	2
DPVols size (GB)	250	15	20	200	2000	380
ASM disk group	DATA	OCR	REDO	TEMP	FRA	N/A
Purpose	OLTP application tablespaces System Sysaux Undo	Oracle Cluster Registry	Online redo Logs Control files	Temp	Incremental backups Archived redo Logs Control file Auto backups	SAN boot OS volumes
Storage Port	1A, 1B, 2A, 2B, 3A, 3B, 4A, 4B					

The following table lists the VSP E990 management server configuration.

Table 6 VSP E990 Configuration for Management Servers

Item	Value/Description
Purpose	VMware Datastores CCI device
RAID level	RAID 10 (2D+2D)
Drive type	3.8 TB NVMe SSD
Number of drives	16
Number of spare drives	1
Number of Pool DEVs	4
LDEV size(s)	1 × 3000 GB
Number and size of CCI devices	1 × 100 MB

Item	Value/Description
Storage port for management servers	7A, 7B, 8A, 8B

4 × RAID 10 groups consisting of 12 × 3.84 TB NVMe SSD drives configured as RAID 10 (2D +2D) were 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 RAID group for the management server cluster, this solution can be configured as dedicated RAID groups, a dedicated HDP pool, or it can use capacity on the HDP pool configured for the Oracle environment depending on customer requirements.

Database layout

The database layout design uses recommended best practices from Hitachi Vantara for Hitachi Virtual Storage Platform 5600 for small random I/O traffic, such as OLTP transactions. The layout design also considers Oracle ASM best practices when using Hitachi storage.

Base the storage design for database layout needs on the requirements of a specific application implementation. 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 to be used in various deployment scenarios to provide the right balance between performance and ease of management for a given scenario.

Oracle configuration

- Data and Indexes Tablespace - Assign an ASM diskgroup with external redundancy for the data and index tablespaces.
- TEMP Tablespace - Place the TEMP tablespace in this configuration in the TEMP ASM disk group.
- Undo Tablespace - Create an UNDO tablespace in this configuration within the Oracle Data ASM disk group. Assign one UNDO tablespace for each node in the Oracle RAC environment.
- Online Redo Logs - Create an ASM diskgroup with external redundancy for Oracle online redo logs.
- Oracle Cluster Registry and Voting Disk - Create an ASM diskgroup with normal redundancy to contain the OCR and voting disks and to protect against single disk failure to avoid loss of cluster availability. Place each of these files in this configuration in the OCR ASM disk groups.

Oracle initial parameters

The following table lists Oracle ASM and database parameters.

Table 7 Oracle ASM and Database Parameters

Category	Item	Value
Oracle RAC option	RAC configuration	Yes
	ASM	Yes – to support Oracle RAC database
Oracle Database environment parameters	SGA_TARGET	384 GB
	PGA_AGGREGATE_TARGET	154 GB
	DB_CACHE_SIZE	172 GB
	DB_KEEP_CACHE_SIZE	76 GB
	DB_RECYCLE_CACHE_SIZE	20 GB
	LOG_BUFFER	512 MB
	USE_LARGE_PAGES	TRUE
	FILESYSTEMIO_OPTIONS	SETALL
	DB_FILE_MULTIBLOCK_READ_COUNT	128
	DISK_ASYNCH_IO	TRUE

Oracle ASM disk mappings

The following table lists volumes, LUNs, and ASM disk groups.

Table 8 Volumes, LUNS, and ASM Disk Groups

ASM Disk Group	ASM Disk	LUN Details	Purpose
OCR	OCR1 - OCR3	3 × 15 GB	Oracle cluster registry and voting disk
REDO	REDO01 - REDO16	32 × 20 GB	Online REDO log group
FRA	FRA01 - FRA16	16 × 2000 GB	Flash recovery area
TEMP	TEMP1 – TEMP16	16 × 200 GB	Temp
DATA	DATA001 – DATA128	128 × 200 GB	Application data

Oracle server configuration

The following table lists the operating system configurations for Oracle servers.

Table 9 Oracle Server OS Configurations

Server Configuration	Server OS Setting Details
RPMs for Oracle Database 19c	bc binutils compat-libcap1 compat-libstdc++ elfutils-libelf elfutils-libelf-devel fontconfig-devel glibc glibc-devel ksh libaio libaio-devel libXrender libXrender-devel libX11 libXau libXi libXtst libgcc libstdc++ libstdc++-devel libxcb make net-tools (for Oracle RAC and Oracle Clusterware) nfs-utils (for Oracle ACFS) python (for Oracle ACFS Remote) python-configshell (for Oracle ACFS Remote) python-rtslib (for Oracle ACFS Remote) python-six (for Oracle ACFS Remote)

Server Configuration	Server OS Setting Details
	targetcli (for Oracle ACFS Remote) openssh smartmontools sysstat oracleasm-support (2.1.12) oracleasm-lib (2.0.17)
/etc/multipath.conf	user_friendly_names: yes find_multipaths: yes path_grouping_policy: multibus path_selector: "service-time 0"
Swap space	64 GB
ORACLEASM_SCANORDER	For the current environment, the value is set to 'mpath' Modify the 'ORACLEASM_SCANORDER' value in the /etc/sysconfig/oracleasm file according to the multipath disk name string in the environment.
Parameter 'path_selector'	For OLTP Database, set "service-time 0" for path_selector in the /etc/multipath.conf file for the best performance.

Management server configuration

The following table lists the VSP E990 configuration for the management server cluster.

Table 10 VSP E990 Management Server Cluster Configuration

Attribute	Value
Purpose	VMware shared datastores Command device
Number and size of LDEVs for datastores	1 × 3000 GB
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 two Hitachi Advanced Server DS220 G2 servers with 3rd Generation Intel Xeon Scalable Processors (Ice Lake) for each storage system architecture that was tested.

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

Two Hitachi Advanced Server DS120 G2 servers are used for VMware ESXi management server configuration.

The following table lists the details of the server configurations for this solution.

Table 11 Hitachi Advanced Server DS220 G2 and DS120 G2 Server Specifications

Hitachi Advanced Server	Server	Host Name	Role	CPU Core	RAM
DS220 G2	Oracle Server 1	sgrac01	Oracle RAC node 1	76 (2 × 38C)	768 GB (32 GB × 24)
	Oracle Server 2	sgrac02	Oracle RAC node 2	76 (2 × 38C)	768 GB (64 GB × 24)
DS120 G2	Management Server 1	VMware ESXi 1	Management VM hosts	24 (2 × 12C)	256 GB (32 GB × 8)
	Management Server 2	VMware ESXi 2		24 (2 × 12C)	256 GB (32 GB × 8)

SAN architecture

Map the provisioned LDEVs to multiple ports on each Hitachi Virtual Storage Platform 5600 (VSP 5600). These LDEV port assignments provide multiple paths to the storage system from the host for high availability. This reference architecture uses two dual port Emulex HBAs for each DS220 G2 server.

- 8 SAN switch connections are used for VSP 5600 Fibre Channel ports.
- 8 SAN switch connections are used for server HBA ports.
- 4 SAN switch connections are used for VSP E990 Fibre Channel ports.
- 4 SAN switch connections are used for management server HBA ports.

The following table lists the Fibre Channel switch connection configuration on the Virtual Storage System and database servers.

Table 12 Fibre Channel Switch Configuration

Server/ Storage System	HBA Ports	Switch Zone	Storage Port	Brocade G620 Switch
DS220 G2 Server 1 VSP 5600	HBA1_1	DS220_G2_62_HBA1_1_VSP_5600_33_1A	1A	SAN-switch 1
	HBA1_2	DS220_G2_62_HBA1_2_VSP_5600_33_2A	2A	SAN-switch 2
	HBA2_1	DS220_G2_62_HBA2_1_VSP_5600_33_1B	1B	SAN-switch 1
	HBA2_2	DS220_G2_62_HBA2_2_VSP_5600_33_2B	2B	SAN-switch 2
DS220 G2 Server 2 VSP 5600	HBA1_1	DS220_G2_64_HBA1_1_VSP_5600_33_1C	1C	SAN-switch 1
	HBA1_2	DS220_G2_64_HBA1_2_VSP_5600_33_2C	2C	SAN-switch 2
	HBA2_1	DS220_G2_64_HBA2_1_VSP_5600_33_1D	1D	SAN-switch 1
	HBA2_2	DS220_G2_64_HBA2_2_VSP_5600_33_2D	2D	SAN-switch 2
DS120 G2 Server 1 VSP E990	HBA1_1	DS120_68_HBA1_1_ASE32_50_7A	7A	SAN-switch 1
	HBA1_2	DS120_68_HBA1_2_ASE32_50_8A	8A	SAN-switch 2
DS120 G2 Server 2 VSP E990	HBA1_1	DS120_69_HBA1_1_ASE32_50_7B	7B	SAN-switch 1
	HBA1_2	DS120_69_HBA1_2_ASE32_50_8B	8B	SAN-switch 2



Note: In a production environment, separate storage ports are recommended for the management servers to avoid impacting database performance. Shared storage ports can be used; however, port utilization should be monitored to avoid performance issues in high-performance environments.

Network architecture

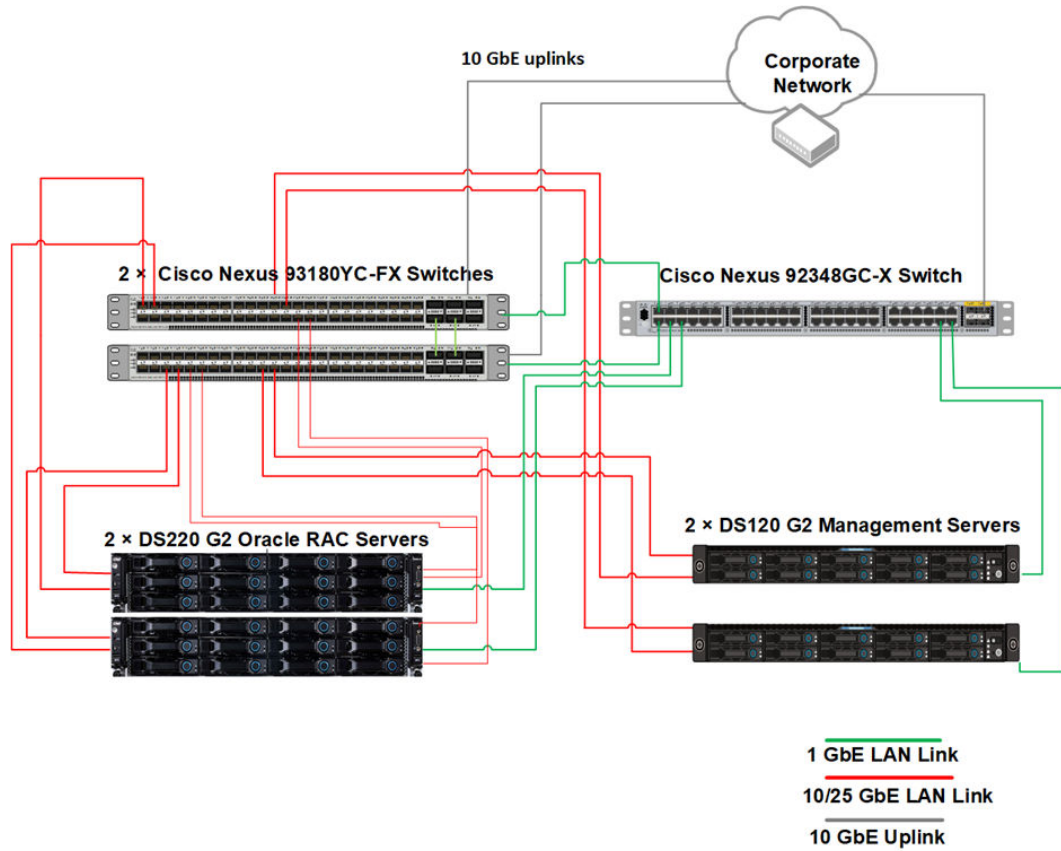
This architecture uses the following separate networks:

- Private Network (also called cluster interconnect) -This network must be scalable. In addition, it must meet the low latency needs of the network traffic generated by the cache synchronization of Oracle Real Application Clusters and inter-node communication among the nodes in the cluster.
- Public Oracle Network -This network provides client connections to Oracle Real Application Clusters and other applications.
- Management Network - This network is for hardware management console connections.

Hitachi Vantara recommends using pairs of 25 Gbps NICs for the cluster interconnect network and public network. Observe these points when configuring private and public networks in your environment:

- For each server in the clusterware configuration, use at least two identical, high-bandwidth, low-latency NICs for the interconnection.
- 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 the database.
- Use at least two private NICs for the cluster interconnection.

The following figure shows the IP network switch connection.



The following table lists the network configuration. Your values might be different.

When creating NIC bonding pairs, ports should be used on different cards to avoid single points of failure (SPoF).

Table 13 Database Server Network Configuration

Server	NIC Port	NIC Bond	IP Address	Network	Bandwidth (Gbps)	Cisco Nexus 93180YC-FX Switch	
						Switch Number	Port
Oracle Database Server1	NIC 1 PORT 1	Bond0	192.168.208.91	Private	25	1	31
	NIC 2 PORT 1				25	2	
	NIC 1 PORT 2	Bond1	192.168.242.91	Public Oracle	25	1	32
	NIC 2 PORT 2				25	2	
	BMC-Dedicated NIC	-	192.168.242.161	Management	1	-	
Oracle Database Server2	NIC 1 PORT 1	Bond0	192.168.208.92	Private	25	1	34
	NIC 2 PORT 1				25	2	
	NIC 1 PORT 2	Bond1	192.168.242.92	Public Oracle	25	1	35
	NIC 2 PORT 2				25	2	
	BMC-Dedicated NIC	-	192.168.242.162	Management	1	-	
Management Server 1	BMC-Dedicated NIC	-	192.168.242.169	Management	1	-	
	NIC 1	-	192.168.244.101	Public Oracle	25	1	

Server	NIC Port	NIC Bond	IP Address	Network	Bandwidth (Gbps)	Cisco Nexus 93180YC-FX Switch	
						Switch Number	Port
	PORT 1						
Management Server 2	BMC-Dedicated NIC	-	192.168.242.170	Management	1	-	
	NIC 1 PORT 1	-	192.168.244.102	Public Oracle	25	1	

The following table lists the virtual IP address and SCAN name configuration that was used when testing the environment.

Table 14 Virtual IP Address and Scan Name Configuration

Server	Virtual IP Address	Scan Name pub-scan
Database Server 1	192.168.242.91	192.168.242.222
Database Server 2	192.168.242.92	192.168.242.223 192.168.242.224

The following table lists the virtual machine configurations running on the management server cluster. Virtual machine configurations listed in this table were 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
vCenter	2	10 GB	300 GB	192.168.242.102
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 detail view	4	10 GB	110 GB	192.168.242.197

Engineering validation

This section summarizes the key observations from the test results for the Hitachi Unified Compute Platform architecture for Oracle RAC deployment with Hitachi Virtual Storage Platform 5600.

Test methodology

The test methodology for this reference architecture used Orion and peakmarks benchmark tools. The peakmarks 9.6 tool was used to validate this solution.

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 at <https://docs.oracle.com/en/database/oracle/oracle-database/12.2/tgdba/IO-configuration-and-design.html#GUID-B89C41F4-B12A-49E1-BFB9-5BF978E2D6B3>.

peakmarks

Test results are demonstrated using peakmarks OLTP test cases.

peakmarks is benchmark software for Oracle platforms. It is used in our tests for the purposes of:

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

This provides transparency and comparability in price versus performance considerations for Oracle infrastructure solutions.

Database configuration

The following table lists parameter details for a two-node Oracle Real Application Clusters ASM database.

Table 16 Two-node Oracle RAC ASM Database Parameters

Oracle Database Parameter	Value
Compatible	19.12.0.0.0
cluster_database	TRUE
cluster_database_instances	2
Oracle Database size	16 TB
Database storage type	ASM
Database fill factor	80%

Test environment

The following table lists the test environment details.

Table 17 Test Environment Details

Item	Description	Value
Operating system	-	OL 8.5
Workload type	-	OLTP/OLAP
Database size	-	16 TB
Number of physical cores	CPU cores	152
Memory	Server memory	768 GB per node
Network	Cluster interconnect	2 × 25 Gbps NIC bonding

Test results

The following table lists Orion test results with VSP 5600 RAID 6 (6D+2P) protection.

Test Case	Test Result	
	Throughput	RT (ms)
8KB Random Read	2,228,200 IOPS	0.44
8KB Random Write	220,039 IOPS	0.56
1MB Sequential Read	24.56 GBps	N/A
1MB Sequential Write	16.49 GBps	N/A

The following table lists peakmarks test results with VSP 5600 RAID 6 (6D+2P) protection.

peakmarks workload	VSP Compression Disabled		Storage Array Based Compression Enabled	
	Throughput	RT (ms)	Throughput	RT (ms)
Storage Performance				
STO-READ SQL sequential read	22.77 GBps	N/A	24.19 GBps	N/A
STO-Random SQL random read via index access, 100% read	2,055,631 IOPS	0.86	1,485,509 IOPS	0.96
STO- Random SQL random read via index access, 80% read	1,835,323 IOPS	0.69	1,384,515 IOPS	0.99
STO- Random SQL random read via index access, 70% read	1,517,165 IOPS	0.79	1,242,037 IOPS	0.75
Server Performance				
SRV-SCAN buffer cache scan rate, full table scan	274.96 GBps	N/A	307.00 GBps	N/A
Log Writer Performance				
LGWR-THR throughput REDO data	2.94 GBps	N/A	3.02 GBps	N/A
DB Writer Performance				
DBWR-THR throughput of database, massive block changes in the buffer cache	9.73 GBps	N/A	7.92 GBps	N/A
Data Analytics Performance				

peakmarks workload	VSP Compression Disabled		Storage Array Based Compression Enabled	
	Throughput	RT (ms)	Throughput	RT (ms)
DA-ROWSTORE data scan row store	236.33 GBps	N/A	247.75 GBps	N/A
DA-STORAGE data scan storage system	24.54 GBps	N/A	20.48 GBps	N/A
DA-COLSTORE data scan column store	3,240.99 GBps	N/A	3248.05 GBps	N/A
Transaction Processing Performance				
TP-LIGHT light transaction, select/update single row via index per transaction 1 rpt 0%update	1,634,737 tps	0.52	1,244,268 tps	0.62
TP-LIGHT light transaction, select/update single row via index per transaction 1rpt 20 %update	403,768 tps	0.95	528,973 tps	1.15
TP-LIGHT light transaction, select/update single row via index per transaction 1rpt 30%update	291,094 tps	0.88	408,415 tps	1.25
TP-MEDIUM transaction, select/update 25 rows via index per transaction 25 rpt 0%update	97,147 tps	3.95	22,861 tps	33.15
TP-MEDIUM transaction, select/update 25 rows via index per transaction rpt 20%update	18,804 tps	6.80	18,464 tps	16.45
TP-MEDIUM transaction, select/update 25 rows via index per transaction 25 rpt 30%update	11,748 tps	10.89	15,466 tps	33.06

peakmarks® test descriptions

The following table lists peakmarks® Key Performance Tests/Metrics for platform components in database operations.

Category	Key Performance Metric	peakmarks® Workload
Server System All accessed data is stored entirely in the database buffer cache. No I/O operations.	Query throughput and response time for simple queries	SRV-QUERY1
	Query throughput and response time for more complex queries	SRV-QUERY25
	Throughput logical reads for online reports	SRV-REPORT
	Scan throughput database buffer cache	SRV-SCAN
	Query throughput and response time for mixed queries and scans	SRV-MIXED
Storage System	SQL sequential I/O throughput	STO-READ
	SQL sequential I/O throughput - using smart scan (offload)	STO-OFFLOAD
	SQL random read throughput and service time - 100% read	STO-RANDOM
	SQL random read throughput and service time - 80% read	STO-RANDOM
	SQL random write throughput	STO-SCATTER

The following table lists peakmarks® Key Performance Metrics for critical database background processes.

Category	Key Performance Metric	peakmarks® Workload
Log Writer (LGWR)	Commit throughput and latency for small transactions	LGWR-LAT1
	Commit throughput and latency for medium-sized transactions	LGWR-LAT25
	Commit throughput and latency for large transactions	LGWR-LAT125
	Log Writer throughput	LGWR-THR
Database Writer (DBWR)	Database Writer throughput	DBWR-THR

The following table lists peakmarks® Key Performance Metrics for representative database operations.

Category	Key Performance Metric	peakmarks® Workload
Data Load	Throughput transactional data load - using the buffer cache	DL-BUFFER
	Throughput data warehouse data load - bypassing the buffer cache	DL-DIRECT
	Throughput IoT data load - using memory-optimized tables	DL-STREAM
Data Analytics	Throughput data scan - using the storage system	DA-STORAGE
	Throughput data scan - using smart scan	DA-OFFLOAD
	Throughput data scan - using row store	DA-ROWSTORE
	Throughput data scan - using column store	DA-COLSTORE
Transaction Processing	Transaction throughput and response time of light transactions	TP-LIGHT
	Transaction throughput and response time of medium transactions	TP-MEDIUM
	Transaction throughput and response time of heavy transactions	TP-HEAVY
	Transaction throughput and response time for mixed transactions	TP-MIXED

The following table lists peakmarks® Key Performance Metrics for PL/SQL application code.

Category	Key Performance Metric	peakmarks® Workload
PL/SQL	Throughput PL/SQL operations	PLS-MIXED
	Execution time in Fibonacci numbers (n = 42), SIMPLE-INTEGERS	PLS-FIBO (SI)

The following is a list of measurement units:

- [s] seconds
- [ms] milliseconds
- [μ s] microseconds
- [qps] queries per second
- [rps] rows per second
- [dbps] database blocks per second
- [tps] transactions per second
- [Mops] million operations per second
- [MBps] megabyte per second
- [IOPS] I/O operations per second

Product descriptions

The following products were used in this solution.

Hitachi Virtual Storage Platform 5600

Hitachi Virtual Storage Platform 5600 supercharges business application performance with all-NVMe storage. It uses Hitachi Ops Center, so you can improve IT operations with the latest AI and ML capabilities. Advanced data reduction in Virtual Storage Platform 5600 enables you to run data reduction with even the most performance hungry applications.

The all-NVMe architecture in Virtual Storage Platform 5600 delivers consistent, low-microsecond latency to reduce latency costs for critical applications. This predictable performance optimizes storage resources.

With Virtual Storage Platform 5600 and the rest of Hitachi's midrange storage family, you have agile and automated data center technology. These systems allow you to cost-effectively meet your current digital expectations and give you the ability to address future challenges, as your application data needs and service levels evolve. With time-tested, proven availability and scalability, Hitachi Vantara delivers infrastructure solutions that help you maximize your data center advantage.

Virtual Storage Platform 5600 supports [Oracle Real Application Clusters](#).

Hitachi Storage Virtualization Operating System RF

[Hitachi Storage Virtualization Operating System RF](#) powers the Hitachi Virtual Storage Platform (VSP) family. 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.

Flash performance is optimized with a patented flash-aware I/O stack, which accelerates data access. Adaptive inline data reduction increases storage efficiency while enabling a balance of data efficiency and application performance. Industry-leading storage virtualization allows SVOS RF to use third-party all-flash and hybrid arrays as storage capacity, consolidating resources for a higher ROI and providing a high-speed front end to slower, less-predictable arrays.

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:

- [Administrator](#): Reduce storage management complexities for the Virtual Storage Platform. Intuitive graphical user interfaces (GUIs) and recommended configuration practices speed storage management operations so you spend more time on strategic efforts, not daily tasks.
- [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 DS220 G2

With a combination of two Intel Xeon Scalable processors and high storage capacity in a 2U rack-space package, [Hitachi Advanced Server DS220 G2](#) delivers the storage and I/O to meet the needs of converged solutions and high-performance applications in the data center.

The Intel Xeon Scalable processor family is optimized to address the growing demands on today's IT infrastructure. The server provides 32 slots for high-speed DDR4 memory, allowing up to 4 TB memory capacity with RDIMM population (128 GB × 32) or 8TB (512 GB × 16) with Intel Optane Persistent Memory population.

DS220 G2 comes in three storage configurations to allow for end user flexibility. The first configuration supports 24 2.5-inch non-volatile memory express (NVMe) drives, the second supports 24 2.5-inch serial-attached SCSI (SAS), serial-ATA (SATA) and up to 8 NVMe drives, and the third supports 12 3.5-inch SAS or SATA and up to 8 NVMe drives. All the configurations support hot-pluggable, front-side-accessible drives as well as 2 optional 2.5-inch rear mounted drives. The DS220 G2 delivers high I/O performance and high capacity for demanding applications and solutions.

Hitachi Advanced Server DS120 G2

With support for two Intel Xeon Scalable processors in just 1U of rack space, the [Hitachi Advanced Server DS120 G2](#) delivers exceptional compute density. It provides flexible memory and storage options to meet the needs of converged and hyperconverged infrastructure solutions, as well as for dedicated application platforms such as internet of things (IoT) and data appliances.

The Intel Xeon Scalable processor family is optimized to address the growing demands on today's IT infrastructure. The server provides 32 slots for high-speed DDR4 memory, allowing up to 4 TB memory capacity with RDIMM population (128 GB × 32) or 8 TB (512 GB × 16) of Intel Optane Persistent Memory. DS120 G2 supports up to 12 hot-pluggable, front-side-accessible 2.5-inch non-volatile memory express (NVMe), serial-attached SCSI (SAS), serial-ATA (SATA) hard disk drive (HDD), or solid-state drives (SSD). The system also offers 2 onboard M.2 slots.

With these options, DS120 G2 can be flexibly configured to address both I/O performance and capacity requirements for a wide range of applications and solutions.

Oracle Linux

[Oracle Linux](#) (OL, formerly known as Oracle Enterprise Linux) is a Linux distribution packaged and freely distributed by Oracle, available partially under the GNU General Public License. It is compiled from Red Hat Enterprise Linux source code, replacing Red Hat branding with Oracle branding.

Oracle Database with Real Application Clusters Option

[Oracle Database](#) has a multi-tenant architecture used to 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 Real Application Clusters](#) (Oracle RAC) is a clustered version of Oracle Database. It is based on a comprehensive high-availability stack that can be used as the foundation of a database cloud system, as well as a shared infrastructure. This ensures high availability, scalability, and agility for any application.

[Oracle Automatic Storage Management](#) (Oracle ASM) is a volume manager and file system for Oracle database files. This supports both 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 ESXi

[VMware ESXi](#) is a foundation for the virtual infrastructure used for the management applications in this architecture. This allows the environment to operate independently from any general-purpose operating system, offering security, reliability, and simplified management.

VMware vCenter Server Appliance

The VMware vCenter Server Appliance is a preconfigured Linux virtual machine, which is optimized for running VMware vCenter Server and the associated services on Linux.

vCenter Server Appliance is an Open Virtualization Format (OVF) template. The appliance is imported to an ESXi host and configured through the web-based interface. It comes pre-installed with all the components needed to run a vCenter Server. These include vCenter SSO (Single Sign-on), Inventory Service, vSphere Web Client, and the vCenter Server itself.

Brocade switches from Broadcom

Brocade and Hitachi Vantara have partnered to deliver storage networking and data center solutions. These solutions reduce complexity and cost, as well as enable virtualization and cloud computing to increase business agility.

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.

Cisco Nexus switches

The Cisco Nexus switch product line provides a series of solutions that make it easier to connect and manage disparate data center resources with software-defined networking (SDN). Leveraging the Cisco Unified Fabric, which unifies storage, data and networking (Ethernet/IP) services, the Nexus switches create an open, programmable network foundation built to support a virtualized data center environment.

Hitachi Vantara

Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA



HitachiVantara.com/contact