

Hitachi Solution for Databases - Oracle Real Application Clusters Database 19c with Hitachi VSP E790

Reference Architecture

Legal Notices

© 2021 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, the Microsoft Corporate 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 at <https://www.hitachivantara.com/en-us/company/legal.html>.

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
Initial release	April 5, 2021

Reference Architecture

This solution uses the high-performing Hitachi Virtual Storage Platform E790 (VSP E790) storage system with NVMe SSDs to boost performance and lower I/O latency. Hitachi Advanced Server DS220 is used in this reference architecture to run a dedicated Oracle Database 19c with the Oracle Real Application Clusters option (Oracle RAC Database), and it uses Oracle Linux 7.9 for the operating system. Hitachi Advanced Server DS120 is dedicated for management applications. This document is intended for:

- 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 E790
- Hitachi Advanced Server DS220
- Hitachi Advanced Server DS120
- Storage area networks
- Oracle Database 19c with RAC
- 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

This solution is engineered, pre-tested, and qualified to provide high performance and high reliability in demanding, dynamic Oracle environments. This reference architecture implements Hitachi Unified Compute Platform for Oracle RAC Database on two nodes using Hitachi Virtual Storage Platform E790. It addresses the high availability, performance, and scalability requirements for OLTP and OLAP workloads. The solution was developed using Hitachi Advanced Server DS220 servers, and DS120 servers with 2nd Generation Intel Xeon Scalable Processors and VSP E990 storage for the management server environment.

Business benefits

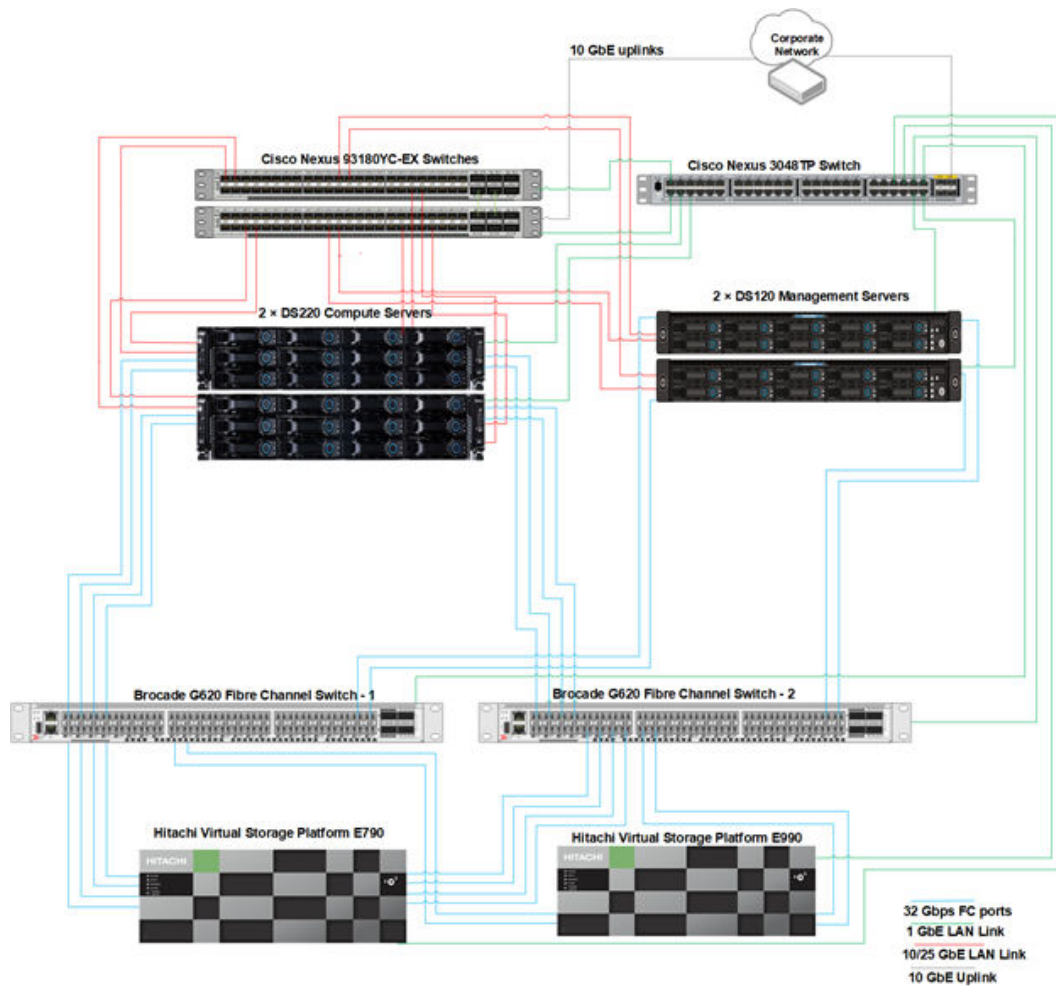
Here are some benefits of this reference architecture:

- Achieves high Oracle Database performance with VSP E790.
- Provides a solution for customers who are looking for very low IO latency for Oracle RAC database.

High-level infrastructure

VSP E790 and Hitachi Advanced Server DS220 were configured with fully redundant hardware for dual fabric connectivity between the hosts and the storage.

The following figure shows the high-level architecture diagram using VSP E790 and Hitachi Advanced Server DS220 for a two-node Oracle 19c RAC configuration, and two Hitachi Advanced Server DS120 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 E790 or other VSP storage arrays.



Note: This solution only requires a single VSP. Two VSPs are shown here because this represents the test lab environment.

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

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

Uplink speed to the corporate network depends on the customer environment and requirements. The Cisco Nexus 93180YC-EX 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. Detailed component information is provided in [Product descriptions \(on page 25\)](#).

Hardware components

Vendor	Hardware	Detail Description	Version	Quantity
Hitachi Vantara	Hitachi Virtual Storage Platform E790	3 × CHB pairs 8 × 32 Gbps Fibre Channel ports 768 GB cache memory 24 × 1.9 TB Backend NVMe SSDs	93-03-21-40/00	1
	Hitachi Virtual Storage Platform E990	2 × controllers 4 × 32 Gbps Fibre Channel ports 16 × 12 Gbps backend SAS ports 1024 GB cache memory 32 × 3.8 TB NVMe SSDs. Larger capacity SSDs can be used.	93_02_02_60/00	1
Hitachi Vantara	Hitachi Advanced Server DS220	2 × Intel Xeon Platinum 8276L 28C CPU @ 2.20GHz 768 GB (64GB × 12) DIMM DDR4 Synchronous Registered (Buffered) 2666 MHz	BIOS: S5BH3B18.H00 BMC: 4.67.06 CPLD: 11	2
		2 × Intel XXV710 dual port 25 GbE NIC cards	Driver: i40e Driver Version: 2.8.20-k	

Vendor	Hardware	Detail Description	Version	Quantity
			Firmware: 7.30	
		2 × Emulex LightPulse LPe32002-M2 2-Port 32 Gb Fibre Channel adapters	Driver: lpfc Driver Version: 12.6.0.3 Firmware: 12.4.243.17	
	Hitachi Advanced Server DS120	2 × Intel Xeon Gold Processor 6240, 18-core, 2.60 GHz, 85W 256 GB (32 GB × 8) DIMM DDR4 Synchronous Registered (Buffered) 2666 MHz 1 × 64 GB MLC SATADOM for boot	BIOS: S5BH3B16.H01 BMC: 4.65.06 CPLD:11	2
		1 × Dual Port 25 GbE NIC Intel XXV710 PCIe card	Driver Version: 1.8.1.6 Firmware: 7.30	
		1 × Emulex LightPulse LPe32002-M2 2-Port 132 Gb Fibre Channel adapters	Firmware: 12.0.193.13 Driver version: 12.0.193.14	
Brocade	G620 Fibre Channel switches	48 × 32 Gbps ports Fibre Channel switch 32 Gbps SFPs	Kernel: 2.6.34.6 Fabric OS : v9.0.0b	2
Cisco	Nexus 93180YC-EX	48 × 10/25 GbE port 6 × 40/100 Gbps Quad SFP (QSFP28) ports	BIOS version: 07.65 NXOS: version 7.0(3)I7(6)	2

Vendor	Hardware	Detail Description	Version	Quantity
	Cisco Nexus 3048TP	48 × 1 GbE 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 might include the SAN, the 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.

SAN and network switches are optional. Direct connect is also possible under certain circumstances, but you should check the support matrix to ensure support prior to implementation. For version details about the hardware components, see the *UCP solutions for Oracle databases and XaaS Interoperability Matrix* at the Hitachi Interoperability Reports site <https://support.hitachivantara.com/en/answers/interoperability.html>.

Software components

Software	Version	Function
Oracle Linux	7.9 with UEK kernel update 5.4.17-2011.6.2.el7u ek.x86_64	Operating system
Oracle Database	19c	Database software
Oracle Real Application Cluster	19c	Cluster 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

Management node software components

Software	Version	Function
ESXi	6.7.0 Build 15160138	Hypervisor for management server
vCenter Server	6.7.0 Build 15976714	VMware cluster management server
Hitachi Device Manager - Storage Navigator Hitachi Storage Advisor Embedded	Microcode-dependent	Storage management software
Hitachi Ops Center	10.5.02	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 E790.

The infrastructure configuration includes the following:

- Oracle RAC servers – Two server nodes are configured in an Oracle Real Application Cluster.
- Management servers – Two servers are used in the management server cluster.
- Storage system – vVols are mapped to each port and presented to the server as LUNs.
- SAN connection – SAN connections from the Fibre Channel HBA ports connect the storage systems through Brocade G620 switches.

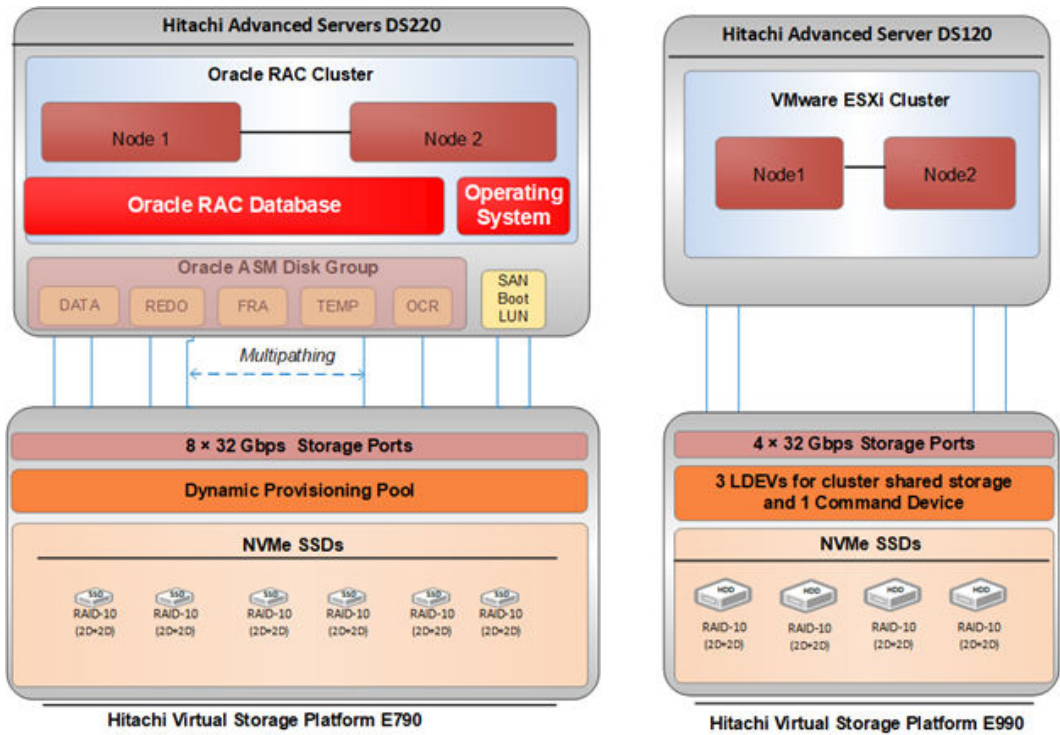
Storage architecture

This section describes the storage architecture for this solution.

Storage configuration

The storage configuration takes into consideration Hitachi Vantara 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.



Note: This solution only requires a single VSP. Two VSPs are shown here because this represents the test lab environment.

This figure lists RAID-10 only. You can configure 3 x PGs with RAID-6 (6D + 2P) protection.

Storage pool configuration

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

Pool ID	Oracle(0)	
Pool Type	Dynamic Provisioning Pool	
RAID Group	1-1 to 1-6	1-1 to 1-3
RAID Level	RAID-10 (2D+2D)	RAID-6 (6D+2P)
Drive Type/Capacity	1.9 TB NVMe SSD	1.9 TB NVMe SSD
Number of Drives	24	24
Number of LDEVs	24	24
LDEV Sizes	880 GB	1320 GB
Pool Capacity	20.61 TB	30.93 TB

Logical storage configuration

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

Dynamic Provisioning Pool ID	Oracle-pool					
Total number of DPVols	128	3	16	8	8	2
DPVols sizes (GB)	120	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	Online redo Logs Control files	Temp	Incremental backups Archived redo Logs Control file Auto backups	SAN boot OS volumes
Storage ports 1A, 1B, 2A, 2B, 3A, 3B, 4A, 4B						

VSP E990 configuration for management servers

The following table lists the VSP E990 configuration for management servers used in this solution.

Item	Value/Description
Purpose	Command Control Interface (CCI) for VMware Datastores
RAID level	RAID-10 (2D+2D)
Drive type	3.8 TB NVMe SSD
Number of drives	24
Number of spare drives	2
Number of Pool DEVs	3
LDEV sizes	1 × 3000 GB
Number and size of CCI devices	1 × 100 MB
Storage ports for management servers	7A, 7B, 8A, 8B

3 × RAID-10 groups consisting of 12 × 3.8 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. The test environment was configured using dedicated RAID groups for the management server cluster. Depending on your requirements, you can configure dedicated RAID groups, a dedicated HDP pool, or use the capacity on the HDP pool configured for the Oracle environment.

Database layout

The database layout design uses recommended best practices from Hitachi Vantara for Hitachi Virtual Storage Platform E790 for small random I/O traffic, such as OLTP transactions. The layout also considers Oracle ASM best practices when using Hitachi Vantara storage.

Specify the storage design for database layout based on the requirements for a specific application implementation. The design can vary from one implementation to another based on the RAID configuration and the 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 the TEMP ASM diskgroup.
- Undo Tablespace - Create an UNDO tablespace within the Oracle Data ASM diskgroup. 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 redologs.
- 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 the OCR ASM diskgroups.

Oracle initial parameters

The following table lists 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_ASYNC_IO	TRUE

Oracle ASM disk mappings

The following table lists ASM disk mapping details.

ASM Disk Group	ASM Disk	LUN Details	Purpose
OCR	OCR1 - OCR3	3 × 15 GB	Oracle cluster registry and voting disk
REDO	REDO01 - REDO16	16 × 20 GB	Online REDO log group
FRA	FRA01 - FRA16	16 × 2000 GB	Flash recovery area
TEMP	TEMP1 – TEMP8	8 × 240 GB	Temp
DATA	DATA001 – DATA128	128 × 200 GB	Application data

Oracle server configuration

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

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-rtlib (for Oracle ACFS Remote) python-six (for Oracle ACFS Remote) targetcli (for Oracle ACFS Remote) openssh

Server Configuration	Server OS Setting Details
	smartmontools sysstat oracleasm-support (2.1.11) oracleasm-lib (2.0.12)
/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
/etc/sysctl.conf	fs.file-max = 6815744 kernel.sem = 250 32000 100 128 kernel.shmmni = 4096 kernel.shmall = 1073741824 kernel.shmmax = 4398046511104 kernel.panic_on_oops = 1 net.core.rmem_default = 262144 net.core.rmem_max = 4194304 net.core.wmem_default = 262144 net.core.wmem_max = 1048576 net.ipv4.conf.all.rp_filter = 2 net.ipv4.conf.default.rp_filter = 2 fs.aio-max-nr = 1048576 net.ipv4.ip_local_port_range = 9000 65500
Swap space	64 GB

Server Configuration	Server OS Setting Details
ORACLEASM_SCANORDER	For the current environment, the value is set to 'mpath' Modify the ORACLEASM_SCANORDER parameter 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 best performance.

Management server configuration

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

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 ports for management servers	7A, 7B, 8A, 8B

Server and application architecture

This reference architecture uses two Hitachi Advanced Server DS220 servers with 2nd Generation Intel Xeon Scalable Processors 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 servers are used to configure VMware ESXi management servers.

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

Hitachi Advanced Server	Server	Host Name	Role	CPU Core	RAM
DS220	Oracle Server 1	sgrac01	Oracle RAC node 1	56 (2 × 28C)	768 GB (64 GB × 12)
	Oracle Server 2	sgrac02	Oracle RAC node 2	56 (2 × 28C)	768 GB (64 GB × 12)
DS120	Management Server 1	VMware ESXi 1	Management VM 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 E790 (VSP E790). These LDEV port assignments provide multiple paths to the storage system from the host for high availability.

This reference architecture uses the following:

- 2 dual port Emulex HBAs per DS220 server.
- 8 SAN switch connections for VSP E790 Fibre Channel ports.
- 8 SAN switch connections for server HBA ports.
- 4 SAN switch connections for VSP E990 Fibre Channel ports.
- 4 SAN switch connections for management server HBA ports.

The following table lists details of the Fibre Channel switch connection configuration on the Hitachi Virtual Storage Platform E790 and database servers.

Server	HBA Ports	Storage Host Group	Switch Zone	Storage Port	Brocade G620 Switch
DS220 Server 1	HBA1_1	DS220_62_HBA1_1	DS220_62_HBA1_1_VSP_E790_33_1A	1A	SAN-switch 1
	HBA1_2	DS220_62_HBA1_2	DS220_62_HBA1_2_VSP_E790_33_2A	2A	SAN-switch 2
	HBA2_1	DS220_62_HBA2_1	DS220_62_HBA2_1_VSP_E790_33_1B	1B	SAN-switch 1
	HBA2_2	DS220_62_HBA2_2	DS220_62_HBA2_2_VSP_E790_33_2B	2B	SAN-switch 2
DS220 Server 2	HBA1_1	DS220_64_HBA1_1	DS220_64_HBA1_1_VSP_E790_33_1C	1C	SAN-switch 1
	HBA1_2	DS220_64_HBA1_2	DS220_64_HBA1_2_VSP_E790_33_2C	2C	SAN-switch 2
	HBA2_1	DS220_64_HBA2_1	DS220_64_HBA2_1_VSP_E790_33_1D	1D	SAN-switch 1
	HBA2_2	DS220_64_HBA2_2	DS220_64_HBA2_2_VSP_E790_33_2D	2D	SAN-switch 2

The following table lists details of the Fibre Channel switch connection configuration on the Hitachi Virtual Storage Platform E990 and database servers.

Server	HBA Ports	Storage Host Group	Switch Zone	Storage Port	Brocade G620 Switch
DS120 Server 1	HBA1_1	DS120_68H BA1_1	DS120_68_HBA1_1_AS E32_50_7A	7A	SAN-switch 1
	HBA1_2	DS120_68_H BA1_2	DS120_68_HBA1_2_AS E32_50_8A	8A	SAN-switch 2
DS120 Server 2	HBA1_1	DS120_69_H BA1	DS120_69_HBA1_1_AS E32_50_7B	7B	SAN-switch 1
	HBA1_2	DS120_69_H BA2	DS120_69_HBA1_2_AS E32_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

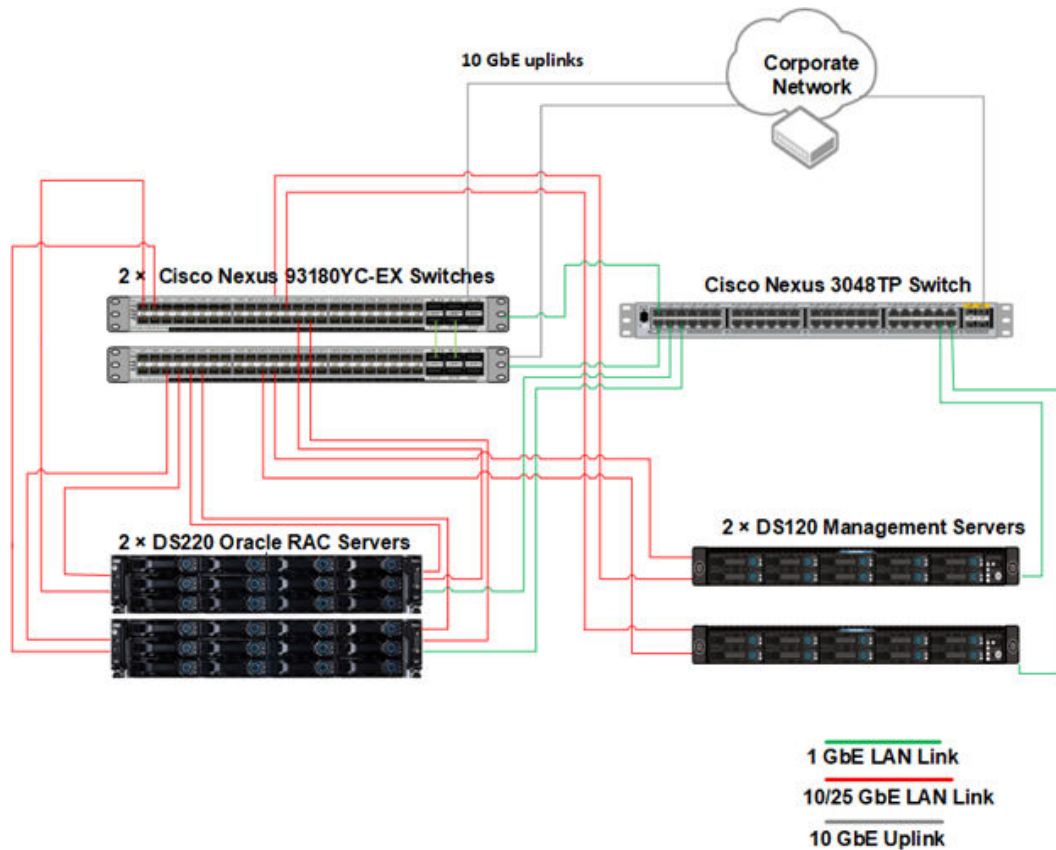
This architecture uses the following networks:

- 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 internode communication 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 database.
- Use at least two private NICs for the cluster interconnection.

The following figure shows the IP network switch connection.



The following tables list the network configuration and the virtual IP address and SCAN name configuration used when testing the environment. Your values might be different.

When creating NIC bonding pairs, ports should be used on different cards to avoid single points of failure (SPoF). The following table lists the database server network configuration.

Ser ver	NIC port	VLAN/ subnet	NIC bon d	IP address	Network	Bandwi dth (Gbps)	Cisco Nexus 93180YC-EX switch	
							Switch number	Port
Ora cle Serv er 1	NIC 1 Port 1	208	Bon d0	192.168.20 8.91	Private	25	1	31
	NIC 2 Port 1					25	2	
	NIC 1 Port 2	242	Bon d1	192.168.24 2.91	Public Oracle	25	1	32
	NIC 2 Port 2					25	2	
	BMC- Dedicate d NIC	242	-	192.168.24 2.161	Manage ment	1		
Ora cle Serv er 2	NIC 1 Port 1	208	Bon d0	192.168.20 8.92	Private	25	1	34
	NIC 2 Port 1					25	2	
	NIC 1 Port 2	242	Bon d1	192.168.24 2.92	Public Oracle	25	1	35
	NIC 2 Port 2					25	2	
	BMC- Dedicate d NIC	242	-	192.168.24 2.162	Manage ment	1		
Mg mt Serv er 1	BMC- Dedicate d NIC	242	-	192.168.24 2.169	Manage ment	1		
	NIC 1 Port 1	244	-	192.168.24 4.101	Public Oracle	25	1	

Server	NIC port	VLAN/ subnet	NIC bond	IP address	Network	Bandwidth (Gbps)	Cisco Nexus 93180YC-EX switch	
							Switch number	Port
Mgmt Server 2	BMC-Dedicated NIC	242	-	192.168.242.170	Management	1	-	
	NIC 1 Port 1	244	-	192.168.244.102	Public Oracle	25	1	

The following table lists the virtual IP addresses 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 are required for the versions used for testing in the lab.

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 Analyzer detail view	4	10 GB	110 GB	192.168.242.197

Engineering validation

This section summarizes the key observations from the test results for this reference architecture.

Test methodology

The test results are demonstrated using Orion I/O workload and peakmarks OLTP test cases. The peakmarks 9.4 tool is 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, see *I/O Configuration and Design* in the [Oracle Database Performance Tuning Guide](#).

peakmarks

[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 (improve efficiency)

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

Database configuration

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

Table 1 Two-node Oracle RAC ASM database parameters

Oracle Database Parameter	Value
Compatible	19.3.0.0.0
cluster_database	TRUE
cluster_database_instances	2
Oracle Database size	8 TB
Database storage type	ASM
Database fill factor	80%

Test environment

The following table lists the test environment details.

Item	Description	Value
Operating system	-	OL 7.9
Workload type	-	OLTP/OLAP
Database size	-	8 TB
Number of physical cores	CPU cores	112
Memory	Server memory	768 GB per node
Network	Cluster interconnect	2 × 25 Gbps NIC bonding

Test results

The following table lists the results of the Orion test with VSP E790 RAID-10 (2D+2D) protection type.

Test Case	Test Result	
8 KB Random Read	Throughput (IOPS or GB/s)	RT (ms)
	2,103,998	0.58
8 KB Random Write	405,210	0.89
1 MB Sequential Read	24.37	N/A
1 MB Sequential Write	4.51	N/A

The following table lists peakmarks results with VSP E790 RAID-10 (2D+2D) and RAID-6 (6D+2P) protection types.

Test Case	RAID-10 (2D+2D) Throughput (IOPS or GB/s)	RT (ms)	RAID-6 (6D+2P) Throughput (IOPS or GB/s)	RT (ms)
STO-READ	25.08	N/A	24.95	N/A
STO-Random 100% read	1,032,157	0.44	754,271	0.47
STO-Random 80% read/20% write	1,080,136	0.51	669,483	0.52
STO-Random 70% read/30% write	696,485	0.58	650,569	0.59
SRV-SCAN	187.72	N/A	162.10	N/A
LGWR-THR	1.91	N/A	2.00	N/A
DBWR-THR	3.77		4.58	
DA-LOW/ROW	181.56		160.44	
DA-LOW/STO	25.40		21.61	
TP-LIGHT 1 rpt 0%update	1,124,721	0.48	980,480	0.46
TP-LIGHT 1rpt 20 %update	221,452	0.99	85,238	1.27
TP-LIGHT 1rpt 30%update	227,687	0.93	83,228	1.30
TP-MEDIUM 25 rpt 0%update	69,346	3.01	56,497	3.69
TP-MEDIUM 25 rpt 20%update	10,512	8.85	2,587	10.74
TP-MEDIUM 25 rpt 30%update	9,433	7.43	2,464	11.19

Product descriptions

These products are used in this reference architecture.

Hitachi Virtual Storage Platform E790

Hitachi Virtual Storage Platform E790 supercharges business application performance with all-NVMe storage. It uses Hitachi Ops Center, so you can improve IT operations with the latest artificial intelligence (AI) and machine learning (ML) capabilities. Advanced data reduction in Virtual Storage Platform E790 enables you to run data reduction with even the most complex applications.

The all-NVMe architecture in Virtual Storage Platform E790 delivers consistent, low-microsecond latency to reduce latency costs for critical applications. This predictable performance optimizes storage resources.

With Virtual Storage Platform E790 and the rest of the Hitachi Vantara midrange storage family, you have agile and automated data center technology. These systems are cost-effective and allow you to cost-effectively meet your current digital expectations. It also gives 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

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:

- 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 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 1U rack-mounted server has the flexibility to power a wide range of solutions and applications.

The highly-scalable memory supports up to 3 TB using 24 slots of high-speed DDR4 memory. Advanced Server 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. This server supports up to 12 small form factor storage devices with up to 4 NVMe drives.

This solution allows you to have a high CPU-to-storage ratio. This is ideal for balanced and compute-heavy workloads.

Multiple CPU and storage devices are available. Contact your Hitachi Vantara sales representative to get the latest list of options.

Hitachi Advanced Server DS220

With a combination of two Intel Xeon Scalable processors and high storage capacity in a 2U rack-space package, [Hitachi Advanced Server DS220](#) 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 24 slots for high-speed DDR4 memory, allowing up to 3 TB of memory per node when 128 GB DIMMs are used. This server supports up to 12 large form factor storage devices and an additional 2 small form factor storage devices.

This server has three storage configuration options:

- 12 large form factor storage devices and an additional 2 small form factor storage devices in the back of the chassis
- 16 SAS or SATA drives, 8 NVMe drives, and an additional 2 small form factor storage devices in the back of the chassis
- 24 SFF devices and an additional 2 SFF storage devices in the back of the chassis

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 Storage Area Network switches

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 | community.HitachiVantara.com

Contact Information

USA: 1-800-446-0744

Global: 1-858-547-4526

HitachiVantara.com/contact