Hitachi Solution for Databases – Non-Virtualized Oracle RAC 19c on HA840 G3 using VSP E790

Reference Architecture Guide

MK-SL-292-00
May 2024

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	May 2024

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 Database using Hitachi Virtual Storage Platform E790.

This Hitachi Solution for Databases with Oracle Real Application Clusters is built on Hitachi Advanced Server HA840 G3 and Virtual Storage Platform E790 storage systems. Hitachi Advanced Server HA840 G3 is a powerful quad core with Intel Xeon 4th generation Scalable Processors that provide reliability, high availability, and scalability while processing all kinds of Oracle workloads.

The VSP E790 storage system with NVMe SSDs is also used in this reference architecture to run a dedicated Oracle Database 19c with the Oracle Real Application Clusters (RAC) on an Oracle Linux operating system. The Hitachi Advanced Server HA810 G3 is used as the management server. In this solution Oracle RAC Database nodes are configured with Oracle Linux 8.6 UEK 6.3 installed.

To configure the solution environment, the following components were used:

- Hitachi Virtual Storage Platform E790 (VSP E790) is used as a storage resource.
- Hitachi Advanced Server HA840 G3 with Intel® Xeon® Platinum 8460H Processor 40C CPUs for storage and computing resources.
- Hitachi Advanced Server HA810 G3 with Intel® Xeon® Silver 4310 Processor 12-core CPUs for the management servers.
- Oracle Real Application Clusters (RAC)

This solution provides the flexibility to select storage and compute resources based on unique requirements. Deploy small databases as well as very large databases, depending on resource availability.

This document is for the following audiences:

- Database administrators
- Storage administrators
- System administrators
- IT professionals responsible for planning and deploying an Oracle Database solution.

To use this document, you need familiarity with the following:

- Hitachi Virtual Storage Platform E790
- Hitachi Advanced Server HA840 G3 servers
- Hitachi Advanced Server HA810 G3 servers
- Storage Area Networks
- Oracle Automatic Storage Management (Oracle ASM)

- Oracle Database 19c
- Oracle Linux

Note: Testing of this configuration was in a lab environment. Many factors 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 E790. It addresses the high availability, performance, and scalability requirements for OLTP and OLAP workloads. This solution was developed on Hitachi Advanced Server HA840 G3, HA810 G3, and VSP E790 storage systems to provide a management server environment.

Business benefits

These are some benefits of this reference architecture:

- Achieve high Oracle Database performance with VSP E790 storage systems.
- Provide a solution for customers who are looking for low I/O latency for Oracle RAC databases.

High-level infrastructure

Hitachi Solution for Databases with Oracle RAC includes the following components:

- Hitachi Advanced Server HA840 G3 servers
- Hitachi Advanced Server HA810 G3 servers
- Hitachi Virtual Storage Platform E790
- Brocade G720 32 Gbps SAN infrastructure
- Cisco 10/25 GbE LAN infrastructure

The configuration of Virtual Storage Platform E790 and Hitachi Advanced Server HA840 G3 has the following characteristics:

- Fully redundant hardware
- Dual fabric connectivity between hosts and storage

Note: While this reference architecture uses the VSP E790, the solution supports the entire Hitachi storage portfolio.

The following figure shows the high-level infrastructure for this solution.

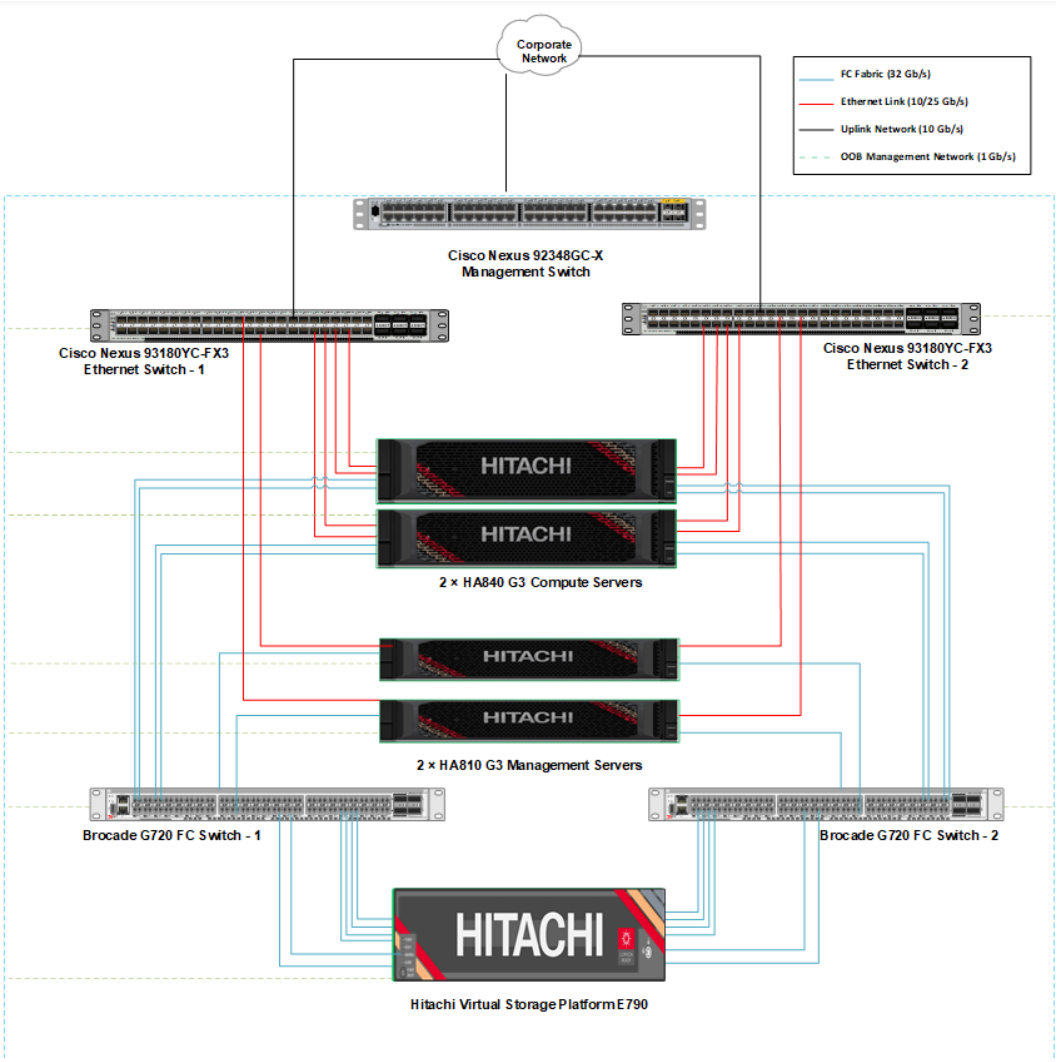

To avoid any performance impact to the production database, Hitachi Vantara recommends using a configuration with the following characteristics:

- A dedicated storage system for the production database
- A dedicated storage system for storing backup data, if needed

The uplink speed to the corporate network depends on the customer environment and requirements. The Cisco Nexus 93180YC-FX3 switches used in this reference architecture can support uplink speeds of 40 GbE or 100 GbE if higher bandwidth is required.

Key solution components

The key solution components for this solution are listed in the following tables. Detailed component information is provided in Product descriptions (on page 19).

Hardware components

The following table lists the hardware components used in this solution.

Vendor	Hardware	Detail Description	Version	Quantity
Hitachi Vantara	Virtual Storage Platform E790	6 × CHA pairs (8 × 32 Gbps Fibre Channel ports in use) 594 GB cache memory 24 x 1.9 TB NVMe SSDs*	93-07-00-40/83	1
Hitachi Vantara	Advanced Server HA840 G3	4 × Intel(R) Xeon(R) Platinum 8460H 40 Cores CPUs @ 2.20GHz 64-bit 4096 GB (64 DIMM × 64 GB) memory RDIMM DDR5-4800 MHz	iLO 6: 1.53 Oct 10 2023 System ROM: U59 v1.46 (09/26/2023)	2
		2 × SN1700E 64GB 2p FC HBA	Firmware: 14.2.589.5 Driver: lpfc Driver version: 12.8.0.11	
		2 × 10/25Gb Dual port SFP28 BCM57414 OCP3 Adapter	Driver: ice Driver version: 0.8.2-k (inbox) Firmware: 223.1.96.0	
		2 × 800 GB NVMe Gen4 Mainstream SFF BC U.3 static SSD	Firmware: 1.24	
Brocade	G720 Fibre Channel switches	64 × 32 Gbps ports Fibre Channel switch 32 Gbps SFPs	Fabric OS: v9.1.1b	2
Cisco	Cisco Nexus C93180YC-FX3	48 × 10/25 GbE ports 6 × 40/100 Gbps Quad SFP (QSFP28) ports	BIOS: version 07.66 NXOS: version 9.3.8	2
	Cisco-C92348GC-X	48 ×1 GbE ports Ethernet switches	BIOS: version 05.42 NXOS: version 10.1(1)	1

Note: The solution was tested with PCIe and OCP Mezzanine NIC cards. Using all PCIe cards is recommended for consistency and better NIC bonding options. SATADOM, SAN boot, or local boot can be used for the boot option. You can choose larger capacity SSDs to fit your business requirements.

Software components

The following table lists software components for compute nodes.

Software	Version	Function
Oracle Linux	8.6 UEK6.3	Operating system for compute nodes
RedHat Linux	8.6 or later	Operating system for compute nodes
Oracle Database	19c (Version 19.21.0.0.0)	Database software
Oracle Real Application Clusters	19c (Version 19.21.0.0.0)	Cluster software
Oracle Grid Infrastructure	19c (Version 19.21.0.0.0)	Volume management, file system software, and Oracle automatic storage management

Note: This solution was tested with RedHat Enterprise Linux 8.6 and later versions.

The following table lists software components for management nodes.

Software	Version	Function
ESXi	8.0 U1	Hypervisor for management server
vCenter Server	8.0 U1	VMware cluster management server
Hitachi Device Manager - Storage Navigator	Microcode dependent	Storage management software
Hitachi Ops Center	10.9.2	Hitachi infrastructure management software

Solution design

This section describes the reference architecture environment to implement Hitachi Unified Compute Platform (UCP) Solution for Oracle with a Real Application Clusters option. The environment uses Hitachi Virtual Storage Platform E790.

The infrastructure configuration includes the following:

- Oracle compute node — There are two hosts configured in an Oracle database environment.
- 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 Virtual Storage Platform and Oracle recommended best practices for the design and deployment of database storage.

The following figure illustrates the high-level storage configuration for this solution.

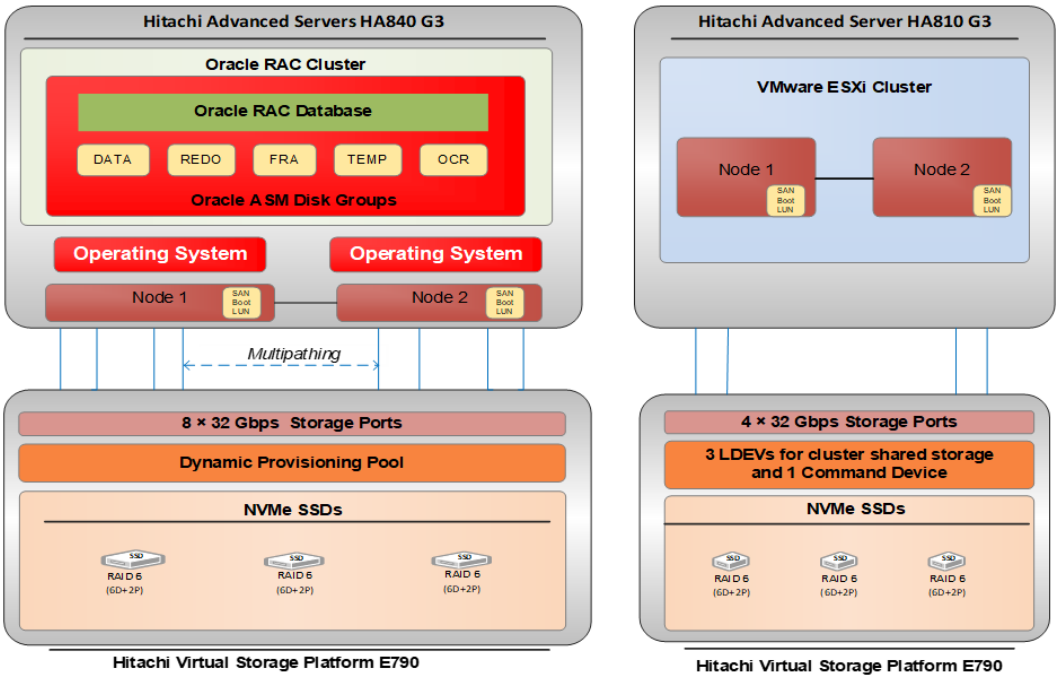

The following table shows the VSP E790 storage pool configuration used for this solution.

Pool ID	Oracle RAC-Pool
Pool Type	Dynamic Provisioning Pool
RAID Group	1-1 to 1-3
RAID Level	RAID 6 (6D+2P)
Drive Type	1.9 TB SSD

Pool ID	Oracle RAC-Pool
Number of Drives	24
Number of LDEVs	24
LDEV Sizes	1320 GB
Pool Capacity	45.6 TB

In this solution, servers use SAN boot with RAID 6 protection. The following table shows the VSP E790 logical storage configuration used in this solution.

ASM disk group	Total number of dynamic provisioning volumes	Dynamic provisioning volume sizes (GB)	Purpose	Storage Ports
DATA	15	200 GB	OLTP Application Tablespaces, System Sysaux, Undo	1A, 2A, 3A, 4A, 5A, 6A, 7A, 8A
REDO	4	20 GB	Online Redo Logs and Control Files	
TEMP	2	200 GB	Temp Tablespace	
FRA	2	200 GB	Incremental Backups, Archived Redo Logs, Control File and Auto backups	
OCR	3	15 GB	Oracle Cluster Registry and Voting Disk	

The following table provides details of the VSP E790 configuration for management servers.

Item	Value/Description
Purpose	VMware datastores CCI device
RAID Level	RAID 6 (6D+2P)
Drive Type	1.9 TB SSD

Item	Value/Description
Number of Drives	24
Number of LDEVs	3
LDEV Size(s)	3 x 1024 GB
Number and Size of CCI device	1 × 100 MB
Storage Port for Management Servers	5C, 6C, 5D, 6D

3× RAID 6 groups consisting of 24 × 1.9 TB NVMe SSD drives configured as RAID 6 (6D +2P) 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 SAS RAID group for the management server cluster, this can be configured as a dedicated SSD RAID group, a dedicated dynamic provisioning pool, and capacity on the dynamic provisioning pool configured for the Oracle environment.

Database layout

The database layout design uses recommended 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.

Base the storage design for database layout needs on the requirements of the specific application implementation. The design can vary greatly from one implementation to another, based on the RAID configuration type and number of drives used in the implementation.

The components in this solution 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 ASM configurations

- Data and Indexes Tablespace — Assign an ASM disk group 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.
- Online Redo Logs — Create an ASM disk group with external redundancy for Oracle online redo logs.
- Oracle Cluster Registry and Voting Disk — Create an ASM disk group 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 shows the Oracle Database settings.

Environment	Value
RAC	Yes
ASM	Yes – to support Oracle RAC Database

Oracle ASM disk mappings

The following table shows the details of the disk mappings from the LUNs to the ASM disk groups for Oracle Database tablespaces for the 2 TB database size. This is an example with a single instance database virtual machine. Adjust parameters accordingly when multiple virtual machine pairs are used.

ASM Disk Group	ASM Disk	UDEV Rules	LUN Details	Purpose
OCR	OCR1-OCR3	/dev/xvd[a-c]1	3 × 15 GB	Oracle cluster registry and voting disk
DATA1	DATA1-DATA15	/dev/xvd[d-r]1	15 × 200 GB	Application data
REDO	REDO01-REDO02	/dev/xvd[s-v]1,	4 × 20 GB	Online REDO log group
FRA	FRA1	/dev/xvd[w-x]1	2 × 200 GB	Flash recovery area
TEMP	TEMP1	/dev/xvd[y-z]1	2 × 200 GB	Temporary Tablespace

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	oracle-database-preinstall-19c
/etc/multipath.conf	user_friendly_names: yes find_multipaths: yes path_grouping_policy: multibus path_selector: "service-time 0”
Swap space	32 GB

Server Configuration	Server OS Setting Details
udev rule	Used 99-oracle-asmdevices.rules file to define device persistency rules
Parameter ‘path_selector’	For OLTP Database, set "`service-time 0`"for path_selection the */etc/multipath.conf*file for the best performance.

Management server configurations

The following table lists management server VM configuration details.

Virtual Machine	vCPU	Virtual Memory	Disk capacity	IP Address
vCenter Server	2	10 GB	300 GB	192.168.242.xx
Oracle Linux Virtualization Manager	4	16 GB	200 GB	192.168.242.xx
Hitachi Ops Center Administrator	4	16 GB	100 GB	192.168.242.xx
Hitachi Ops Center Analyzer	4	32 GB	800 GB	192.168.242.xx
Hitachi Ops Center Analyzer detail view	4	10 GB	110 GB	192.168.242.xx

Server and application architecture

This reference architecture uses two Hitachi Advanced Server HA840 G3 servers with 4th Generation Intel Xeon Scalable Processors for each storage system architecture that was tested. Two Hitachi Advanced Server HA810 G3 servers are used for VMware ESXi management server configuration.

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

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

Server Make and Model	Server Host Name	Role	CPU Type	CPU Core	RAM
Hitachi Advanced Server HA840 G3	Compute node 1	Oracle Linux 8.6 RAC node 1	4 × Intel(R) Xeon(R) Platinum 8460H 40 Cores CPUs @2.20GHz 64-bit	160 (4 × 40)	4096 GB (64 GB × 64)

Server Make and Model	Server Host Name	Role	CPU Type	CPU Core	RAM
	Compute node 2	Oracle Linux 8.6 RAC node 2		160 (4 x 40)	4096 GB (64 GB × 64)
Hitachi Advanced Server HA810 G3	ESXi_101	Management server for Hitachi management Applications	2 × Intel Xeon Processor 4310, 12-core	24 (2 × 12C)	256 GB (32 GB × 8)
	ESXi_103			24 (2 × 12C)	256 GB (32 GB × 8)

SAN architecture

Map the provisioned LDEVs to multiple ports on Hitachi Virtual Storage Platform E590H. These LDEV port assignments provide multiple paths to the storage system from the host for high availability. This reference architecture uses four dual port SN1700E 64GB 2p Fibre Channel HBAs per Advanced Server HA840 G3.

Compute servers

- 8 SAN switch connections are used for VSP E790 Fibre Channel ports.
- 8 SAN switch connections are used for server HBA ports.

Management servers

- 4 SAN switch connections are used for VSP E790 Fibre Channel ports.
- 4 SAN switch connections are used for management server HBA ports.

The following table lists details of the Fibre Channel switch connect configuration on VSP E790 ports.

Server	HBA	Host Group Name	Host Name	Switch Zone	Storage System	Storage Port	Brocade G720 Switch
HA840 G3 Server1	HBA1	CN31	CN31_HBA1_1	CN31_HBA1_1_ASE42_43_1A	VSP E790	1A	SW1-port1
	HBA2	CN31	CN31_HBA1_2	CN31_HBA1_2_ASE42_43_2A		2A	SW2-port1
	HBA3	CN31	CN31_HBA2_1	CN31_HBA2_1_ASE42_43_3A		3A	SW1-port2
	HBA4	CN31	CN31_HBA2_2	CN31_HBA2_2_ASE42_43_4A		4A	SW2-port2

Server	HBA	Host Group Name	Host Name	Switch Zone	Storage System	Storage Port	Brocade G720 Switch
HA840 G3 Server2	HBA1	CN32	CN32_HBA1_1	CN32_HBA1_1_ASE42_43_5A		5A	SW1-port3
	HBA2	CN32	CN32_HBA1_2	CN32_HBA1_2_ASE42_43_6A		6A	SW2-port3
	HBA3	CN32	CN32_HBA2_1	CN32_HBA2_1_ASE42_43_7A		7A	SW1-port4
	HBA4	CN32	CN32_HBA2_2	CN32_HBA2_2_ASE42_43_8A		8A	SW2-port4

The following table lists details of the Fibre Channel switch connect configuration on VSP E790 ports for management servers.

Server	HBA	Host Group Name	Host Name	Switch Zone	Storage Port	Brocade G720 Switch
HA810 G3 Server1	HBA1	MN33	MN33_HBA1	MN33_HBA1_ASE42_43_5C	5C	69
	HBA2	MN33	MN33_HBA2	MN33_HBA2_ASE42_43_6C	6C	70
HA810 G3 Server2	HBA1	MN34	MN34_HBA1	MN34_HBA1_ASE42_43_7C	7C	69
	HBA2	MN34	MN34_HBA2	MN34_HBA2_ASE42_43_8C	8C	70

Note: In a production environment, use separate storage ports for the management servers to avoid impact on 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.

Note that it is best practice to use pairs of 25 Gbps NICs for the cluster interconnect network and public network. Observe these guidelines 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 interconnect.

Physical network configuration

The following figure shows the IP network switch connection.

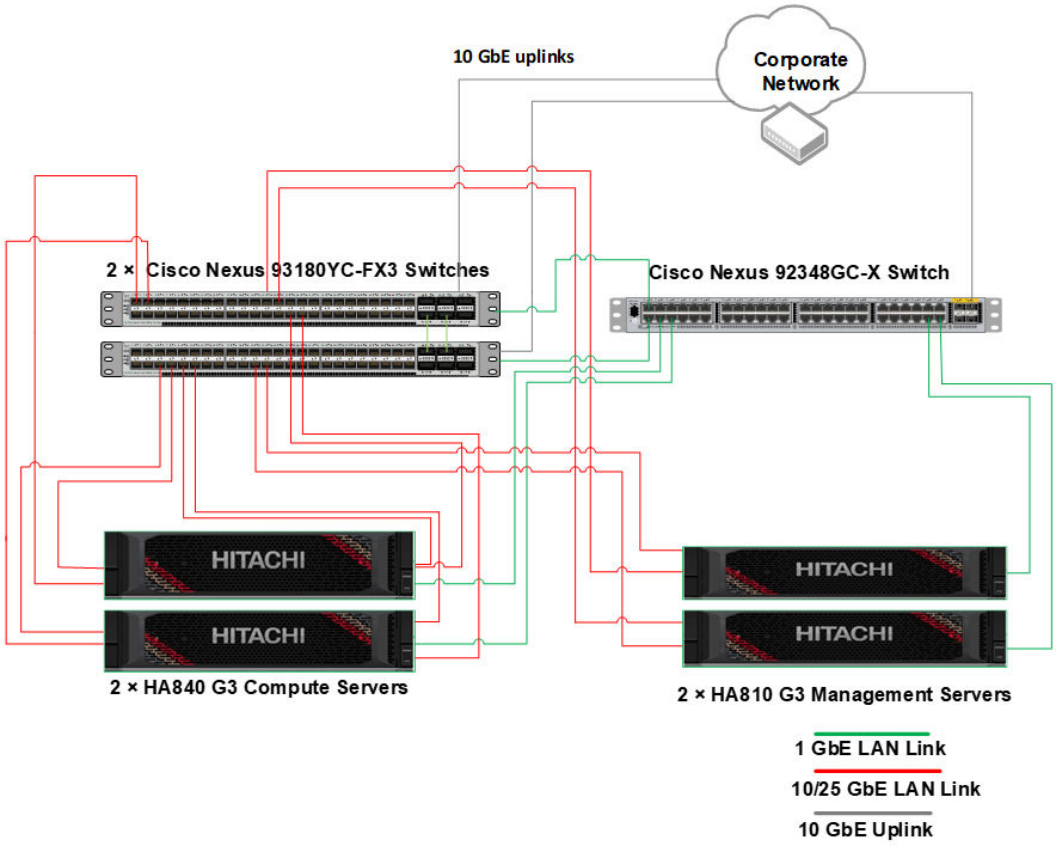

The following table lists the HA840 G3 and HA810 G3 network configuration for this solution.

Server	NIC Ports	Sub-net	NIC Bond	IP Address	Net-work	Band-width (Gbps)	Cisco Nexus 93180YC-FX3 Switch	
							Switch Number	Port
HA840 G3 Server 1	NIC 1-Port 0	242	Bond0	192.198.242.204	Public	25	1	41
	NIC 1-Port 1					25	2	
	NIC 2-Port 0	100	Bond1	192.168.100.xx	Private	25	1	42
	NIC 2-Port 1					25	2	
	iLO-Dedi-cated NIC	242	-	192.168.242.xx	Mgmt	1		
HA840 G3 Server 2	NIC 1-Port 0	242	Bond0	192.198.242.206	Public	25	1	43
	NIC 1-Port 1					25	2	
	NIC 2-Port 0	100	Bond1	192.168.100.xx	Private	25	1	44
	NIC 2-Port 1					25	2	
	iLO-Dedi-cated NIC	242	-	192.168.242.xx	Mgmt	1		
HA810 G3 Mgmt Server 1	NIC 1-Port 0	242	Bond0	192.198.242.101	Public	25	1	49
	NIC 1-Port 1					25	2	
	iLO-Dedi-cated NIC	242	-	192.168.242.xx	Mgmt	1		

Server	NIC Ports	Sub-net	NIC Bond	IP Address	Net-work	Band-width (Gbps)	Cisco Nexus 93180YC-FX3 Switch	
							Switch Number	Port
HA810 G3 Mgmt Server 2	NIC 1-Port 0	242	Bond0	192.198.242.102	Public	25	1	50
	NIC 1-Port 1					25	2	
	iLO-Dedi-cated NIC	242	-	192.168.242.xx	Mgmt	1		

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

The following table lists the network configuration for servers and VSP E790.

Name	IP Address
Oracle Linux host 1	192.168.242.xx
Oracle Linux host 2	192.168.242.xx
Management Server 1	192.168.242.xx
Management Server 2	192.168.242.xx
VSP E790	192.168.242.xx
VSP E790 CTL1	192.168.242.x
VSP E790 CTL2	192.168.242.x

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

Server	Virtual IP Address	Scan Name pub-scan
Database Server 1	192.168.242.xx	192.168.242.xxx
Database Server 2	192.168.242.xx	192.168.242.xxx

Server	Virtual IP Address	Scan Name pub-scan
		192.168.242.xxx

The following table lists the network configuration for the switches in this solution.

Switch Type	Model	Switch Name	IP Address for MGMT port
Cisco 1GbE Management Network Switch	Cisco Nexus C92348GC-X	C92348GC-X -1	192.168.242.xx
Cisco 10G/25GbE Network Switch	Cisco Nexus N9K-C93180YC-FX3	Cisco C93180YC-FX3-1	192.168.242.xx
Cisco 10G/25GbE Network Switch	Cisco Nexus N9K-C93180YC-FX3	Cisco C93180YC-FX3-2	192.168.242.xx
Brocade Fibre Channel SAN Switch	G720	SAN-switch 1	192.168.242.xx
Brocade Fibre Channel SAN Switch	G720	SAN-switch 2	192.168.242.xx

Engineering validation

This section summarizes the key lab verification tests performed on Hitachi Solution for Databases - Oracle Real Application Clusters (RAC) using Hitachi Advanced Server HA840 G3 and Hitachi Virtual Storage Platform E790.

Database configuration

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

Oracle Database Parameter	Value
Compatible	19.21.0.0.0
Oracle Database size	2 TB
Database storage type	ASM
Database fill factor	80%

Test environment

The following table lists configuration details for VSP E790 testing.

Item	Value
Operating System	OL 8.6 UEK6.3
Workload Type	OLTP/OLAP
Database Size	2 TB
Number of CPUs	160 cores per node
Host Cluster Network	2 × 25 Gbps NIC Bonding

Conclusion

Hitachi Advanced Server HA840 G3 and Hitachi Virtual Storage Platform E790 storage system have been tested to provide power to compute nodes in non-virtualized environments to perform heavy transactions and workloads in diverse conditions. At the same time server memory CPU resources were stable.

This solution was also tested by running Oracle RAC database with multiple workloads in parallel for longer durations, and we received good results for OLTP/OLAP transactions. Therefore, Hitachi Vantara hardware provides robust and high-performance infrastructure components, including storage systems, processors, and networking equipment which ensures that the database operates on a reliable and efficient foundation, delivering consistent performance even under heavy workloads.

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 Advanced Server HA840 G3

The Hitachi Advanced Server HA840 G3 is a high-density, four-socket (4S) server with high performance, scalability, and reliability, all in a 2U chassis. Supporting the latest 4th Generation Intel Xeon Scalable Processors, the Hitachi Advanced Server HA840 G3 offers greater processing power, up to 16 TB of DDR5 memory, IO up to six PCIe Gen 5 slots, 2 OCP slots, plus the intelligence and simplicity of automated management with iLO 6. The Hitachi Advanced Server HA840 G3 is the ideal server for business-critical workloads, in-memory databases, data analytics, virtualization, server consolidation, business processing, and general 4S data-intensive applications where data center space and the right performance are paramount.

Hitachi Advanced Server HA810 G3

Optimized for performance, high density, and power efficiency in a dual-processor server, Hitachi Advanced Server HA810 G3 delivers a balance of compute and storage capacity. This rack mounted server has the flexibility to power a wide range of solutions and applications

Highly scalable memory supports up to 4 TB RAM using 32 slots of 2200 MHz DDR5 RDIMM. HA810 G3 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 in this architecture and would be used in most cases:

- vCenter Server
- Hitachi Ops Center
- Oracle Enterprise Manager (OEM) 13c
- Hitachi Storage Adapter for Oracle Enterprise Manager
- Hitachi Server Adapter for Oracle Enterprise Manager
- Oracle Adapter Manager

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

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.

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 with seventh generation 64Gb/sec Fibre Channel interfaces, simplifying scale-out network architectures. Get the high-performance, availability, ease of management, and support for the next generation of Hitachi Virtual Storage Platform storage systems on a solid storage network foundation that can grow as your need grows.

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.

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 is a list of measurement units:

- [s] seconds
- [ms] milliseconds
- [µs] 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

Hitachi Vantara

Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA

HitachiVantara.com/contact