

Hitachi CI Databases Solution — Oracle RAC Database using Hitachi VSP storage and Eviden BullSequana Servers

Reference Architecture Guide

© 2025 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., Hitachi Vantara, Ltd., or Hitachi Vantara Corporation (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, DB2, DS6000, DS8000, Enterprise Storage Server, eServer, FICON, FlashCopy, GDPS, HyperSwap, IBM, OS/390, PowerHA, PowerPC, S/390, System z9, System z10, Tivoli, z/OS, z9, z10, z13, z14, z15, z16, 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>.

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.

Thank you!

Revision history

Changes	Date
Initial release	January 2025

Reference Architecture Guide

This reference architecture guide has been developed by Hitachi Converged Infrastructure (CI) to implement the Hitachi Solution for Databases on Hitachi VSP storage and Eviden BullSequana series (SH20/SH40/SH80) as a compute node equipped with Intel 4th Gen Xeon Scalable Processors which provides a high-performance, low latency, integrated, converged solution for Oracle Database. With these products, you can design an Oracle converged infrastructure to meet your requirements and budget.

This solution uses the high-performing Eviden BullSequana series (SH20/SH40/SH80) servers to boost performance and lower I/O latency. The VSP E1090 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) option on the Oracle Linux 9.1 UEK7U1 operating system. Hitachi Advanced Server HA810 G3/Eviden BullSequana SH20 with Intel Xeon Gold/silver processors is dedicated for management applications.

This document is intended for:

- Database administrators
- Storage administrators
- Database performance analyzers
- IT professionals responsible for planning and deploying an Oracle Database solution

To use this reference architecture guide, you must be familiar with the following:

- Hitachi Virtual Storage Platform E1090
- Eviden BullSequana series SH20/SH40/SH80
- Hitachi Advanced Server HA810 G3
- Storage area networks
- Oracle 19c RAC Database
- Oracle Automatic Storage Management (Oracle ASM)
- Oracle Linux



Note: This configuration was tested in a lab environment. Many factors affect production environments beyond prediction or duplication in a lab environment. Before deployment of this solution, 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 and workload.

Solution overview

The Hitachi solution for Oracle databases on Hitachi Converged Infrastructure (CI) for Oracle RAC 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 Converged Infrastructure (CI) for Oracle Real Application Clusters on two nodes using Hitachi Virtual Storage Platform E1090 storage. It addresses the high availability, performance, and scalability requirements for OLTP and OLAP workloads. This solution was developed using Intel Sapphire Rapids Platform-based BullSequana SH20, HA810 G3, and VSP E1090 storage systems for the management server environment.



Note: Customers can choose Eviden BullSequana series (SH20/SH40/SH80) for compute node according to the scalability and performance requirements of their environment.

Business benefits

Here are some benefits of this reference architecture:

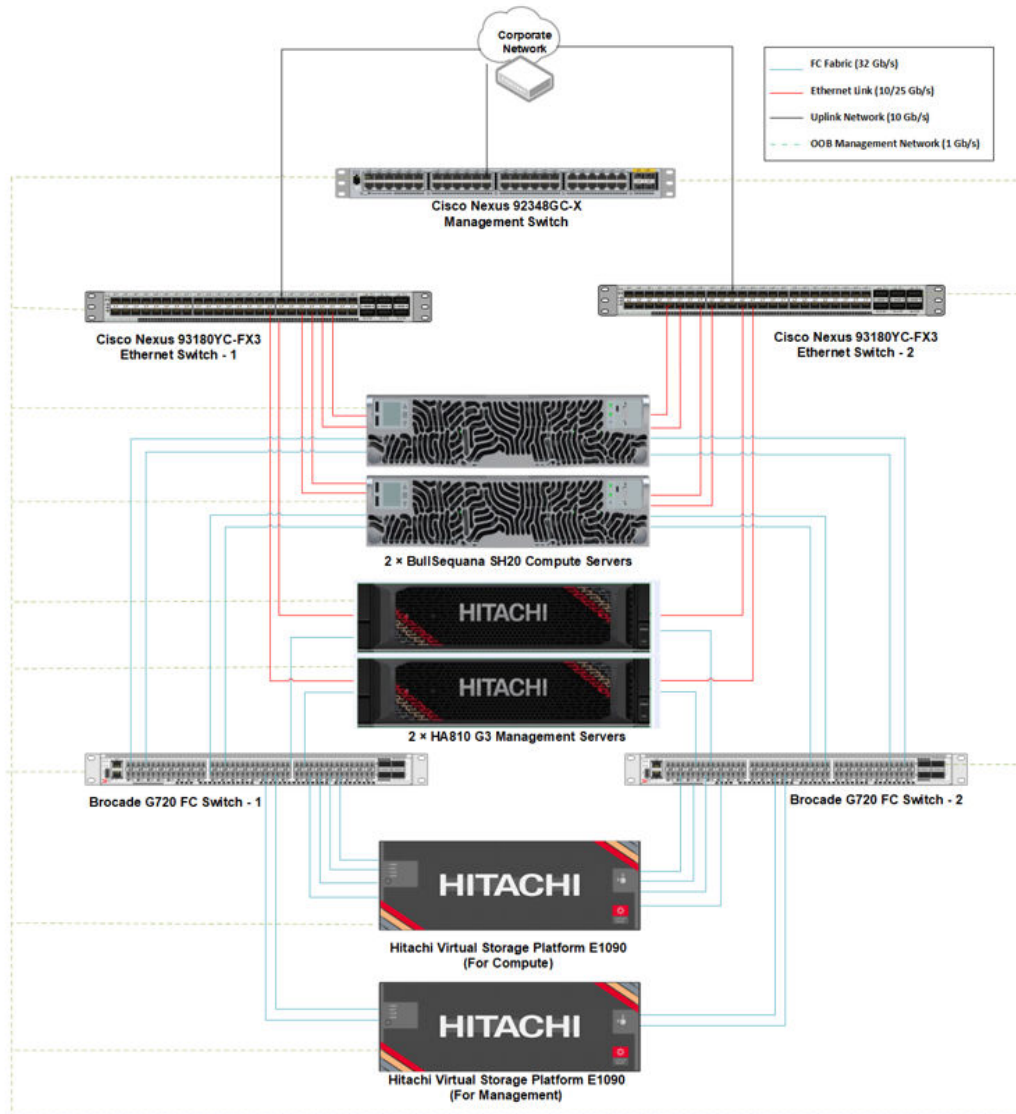
- Achieves best Oracle Database performance with VSP E1090.
- Provides a solution for customers who are looking for highly scalable, redundant infrastructure that can deliver very low I/O latency for Oracle RAC databases.

High level infrastructure

Hitachi VSP E1090 and BullSequana SH20 server 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 E1090 and BullSequana SH20 for a two-node Oracle 19c RAC database configuration, and two HA810 G3 servers for management of the architecture.

Data volumes for the management servers were configured on a VSP E1090, which was shared across different management nodes. The VSP E1090 is not required in your environment. Data volumes can also be configured on the same VSP E1090 or any other VSP storage systems.



Note: Customers can also use Eviden BullSequana series (SH20/SH40/SH80) with Intel Xeon gold/silver processors instead of Hitachi Advanced server HA810 G3 for management appliance.

To avoid a performance impact on 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-FX3 switches can support uplink speeds of 40 GbE or 100 GbE if higher bandwidth is required.

Memory configurations

The BullSequana SH20, SH40, and SH80 series servers offer memory capacities ranging from 512 GB to 16 TB per node, along with external storage. Each 2S node can support up to 32 memory slots which support various DIMM sizes such as 32 GB, 64 GB, and 128 GB.

CPU configurations

The BullSequana SH20, SH40, and SH80 series servers support a wide variety of 4th and 5th Generation Intel Xeon® Scalable processors with innovative capabilities to address the business needs according to their requirements.

Oracle database clusterware can be configured using BullSequana servers with the following CPU models according to the performance and scalability requirements, aligning with the availability of Oracle database licenses.

The BullSequana series servers are capable of scaling from 2 to 32 sockets in a single server with scale in 2 socket increments. Configuring DB appliances for optimal performance, customers can select platinum series CPU models from three distinct categories as follows.

- Database environments requiring the highest number of cores can opt 8490H CPU with 60 cores and 1.9 GHz.
- Database environments requiring a medium number of cores can use 8450H CPU with 28 Cores and 2 GHz.
- Database environments requiring a low number of cores with HIGH frequency can select 8444H CPU with 16 cores 2.9GHz.

Customers can also choose gold CPUs such as G6448H 32 cores 2 GHz , G6416 18 cores 2 GHz or 6434H 8 cores 4 GHz according to performance requirements for the 2 and 4 socket Eviden servers where less expensive and higher frequency CPUs are required. Customers can also save on Oracle licensing costs based on CPU cores by selecting processors with lower cores and high frequency.

CPU Sapphire Rapids			
	2 sockets Only	2S upto 4S	2S - 4S upto 8S
2S	8480+ 56C@2 GHz 8470 52C@2 GHz 8468V 48C@2.4 GHz 8458P 44C@2.7GHz 6444Y 16C@3.6 GHz 6434 8C@3.7 GHz		
4S	6448H 32C@2.4 GHz 6418H 24C@2.1 GHz 6416H 18C@2.2 GHz 6434H 8C@3.7 GHz		
8S	8490H 60C@1.9GHz 8468H 48C@2.1GHz 8460H 40C@2.2GHz 8454H 32C@2.1 GHz 8450H 28C@2GHz 8444H 16C@2.9 GHz		

Key solution components

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

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

Vendor	Hardware	Detailed Description	Version	Quantity
Hitachi Vantara	VSP E1090	2 × CHB pairs 24 × 32 Gbps Fibre Channel ports 768 GB cache memory 24 × 7.6 TB NVMe SSDs	93-07-21-80/00	1

Vendor	Hardware	Detailed Description	Version	Quantity
Eviden	BullSequana SH20	2 × Intel Xeon Platinum 8490H 60C CPUs @ 3.5 GHz 2048 GB (32 GB × 64) RDIMM DDR5 Synchronous Registered (buffered) 4800 MHz	BIOS: BIOS_SAR120.79.02.011 BMC: 152.05.0022	2
		2 × Dual Port 25 GbE Mellanox MT2894 Network PCIe cards	Driver: mlx5_core Driver version: 5.0-0 (inbox) Firmware: 26.35.1012	
		2 × Dual Port 32 GbE NIC Emulex LPe35002-M2 FC HBAs	Driver: lpfc Driver version: 12.8.0.11 Firmware: 14.0.505.12	
	HA810 G3	2 × Intel Xeon processors 4310, 12-core, 2.10 GHz, 120W 256 GB (32 GB × 8) DIMM DDR5-4800 Synchronous Registered (Buffered) 4800 MHz 1 × 480 GB NVMe disk NS204i-u Gen11 Boot Controller	System ROM: U54 v1.30(03/01/2023) iLO6: 1.53 Oct 10 2023	2
		1 × Dual Port 25 GbE NIC Intel E810 PCIe card	Driver: ice Driver Version: 0.8.2-k (inbox) Firmware: 4.10	
		2 × SN1610E 32Gb 2p Fibre Channel HBAs	Driver: lpfc Driver version: 14.0.0.1 Firmware: 14.0.499.29	

Vendor	Hardware	Detailed Description	Version	Quantity
Brocade	G720 Fibre Channel switches	48 × 32 Gbps ports Fibre Channel switches 32 Gbps SFPs	Fabric OS: v9.1.1c	2
Cisco Systems	Nexus 93180YC-FX3	48 × 10/25 GbE ports 6 × 40/100 Gbps Quad SFP (QSFP28) ports	NXOS: version 10.2(5)	2
	Cisco Nexus 92348GC-X	48 × 1 GbE ports Ethernet switches	NXOS: version 10.2(5)	1

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

Certain components are 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.

The following table lists software components used in this solution.

Software	Version	Function
Oracle Linux	9.1 UEK7U1	Operating system
Oracle Database	19.19.0.0.0	Database software
Oracle Real Application Cluster	19.19.0.0.0	Cluster software
Oracle Grid Infrastructure	19.19.0.0.0	Volume management, file system software, and Oracle Automatic Storage Management (ASM)
Oracle ASMLib	2.1.12	Oracle software for ASM

The following table lists the management node software components.

Software	Version	Function
ESXi	7.0 U3	Hypervisor for management server
vCenter Server	7.0 U3	VMware cluster management server

Software	Version	Function
Hitachi Device Manager - Storage Navigator	Microcode dependent	Storage management software
Hitachi Ops Center	10.9.2-00	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 a single Hitachi Virtual Storage Platform E1090.

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 servers as LUNs.
- SAN Connection – There are SAN connections to connect the Fibre Channel HBA ports to the storage through Brocade G720 switches.

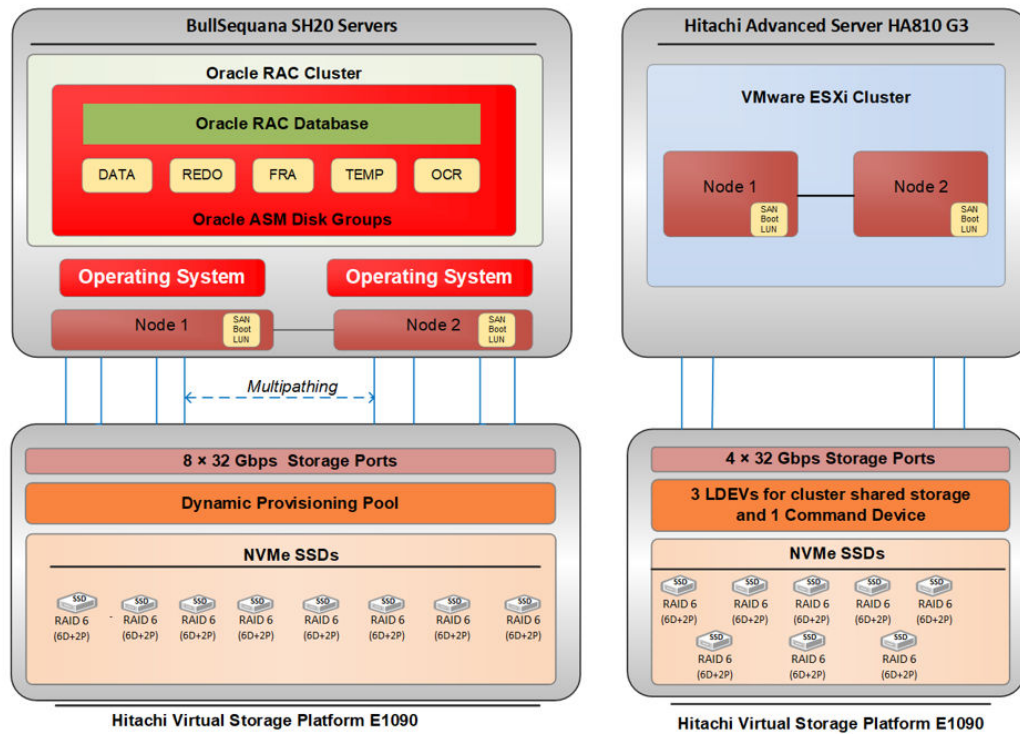
Storage architecture

This section describes the storage architecture for this solution.

Storage configuration

The storage configuration takes into consideration Hitachi Vantara 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 following figure.



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
RAID Protection Type	RAID 6 (6D+2P)
Drive Type/Capacity	7.6 TB NVMe SSD
Number of Drives	16
Number of LDEVs	16
LDEV Size(s)	2560 GB
Pool Capacity	40 TB

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	16	16	2
DPVols size (GB)	200	15	20	200	2000	380
ASM diskgroup	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	3A, 3B, 7A, 7B, 4A, 4B, 8A, 8B					

The following table lists the VSP E1090 management server configuration.

Item	Value/Description
Purpose	VMware Datastores CCI device
RAID level	RAID 6 (6D+2P)
Drive type	1.9 TB NVMe SSD
Number of drives	48
Number of Pool DEVs	4
LDEV size(s)	3 × 1000 GB
Number and size of CCI devices	1 × 100 MB
Storage port for management servers	7A, 7B, 8A, 8B

6 × RAID 6 groups consisting of 48 × 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 necessary. The test environment was configured using a dedicated RAID group for the management cluster. You can also configure this solution using dedicated RAID groups, a dedicated HDP pool, or capacity on the HDP pool configured for the Oracle environment according to customer requirements.

Database layout

The database layout design uses recommended best practices from Hitachi Vantara for Hitachi Virtual Storage Platform E1090 for small random I/O traffic, such as OLTP transactions. The layout design also considers Oracle ASM best practices when using Hitachi storage.

The storage design for the database layout should be based on application-specific requirements. Design can vary greatly across implementations based on RAID configuration and number of drives allocated. The components in this solution have the flexibility to accommodate various deployment scenarios by providing the right balance between performance and ease of management.

Oracle configuration

- Data and Index Tablespaces - Assign an ASM diskgroup with external redundancy for data and index tablespaces.
- TEMP Tablespace - Place the TEMP tablespace in the TEMP ASM diskgroup.
- Undo Tablespace - Create an UNDO tablespace in 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 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 and 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	522 GB
	PGA_AGGREGATE_TARGET	348 GB
	DB_CACHE_SIZE	172 GB
	DB_KEEP_CACHE_SIZE	104 GB
	DB_RECYCLE_CACHE_SIZE	51 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 diskgroups.

ASM Diskgroup	ASM Disk	LUN Details	Purpose
OCR	OCR1 - OCR3	3 × 20 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 – 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.

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	64 GB
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_selector in the /etc/multipath.conf file for the best performance.

Management server configuration

The following table lists the VSP E1090 configuration for the management server cluster.

Attribute	Value
Purpose	VMware shared datastores Command device
Number and size of LDEVs for datastores	3 × 1000 GB
Number and size of command devices	1 × 100 MB
Storage port for management servers	1A, 1B, 2A, 2B

Server and application architecture

This reference architecture uses two BullSequana SH20 servers with 4th Generation Intel Xeon Platinum Processors for each storage system architecture that was tested. Two Hitachi Advanced Server HA810 G3 servers were 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.

Hitachi Advanced Server	Server	Host Name	Role	CPU Core	RAM
BullSequana SH20	Oracle Server 1	sgrac01	Oracle RAC node 1	120 (2 × 60C)	2048 GB (32 GB × 64)
	Oracle Server 2	sgrac02	Oracle RAC node 2	120 (2 × 60C)	2048 GB (32 GB × 64)
HA810 G3	Management Server 1	VMware ESXi 1	Management VM hosts	64 (2 × 32C)	256 GB (16 GB × 16)
	Management Server 2	VMware ESXi 2		64 (2 × 32C)	256 GB (16 GB × 16)

SAN architecture

Map the provisioned LDEVs to multiple ports on each Hitachi Virtual Storage Platform E1090 (VSP E1090). 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 BullSequana SH20 server.

The SAN architecture includes:

- 8 SAN switch connections for VSP E1090 Fibre Channel ports.
- 8 SAN switch connections for server HBA ports.
- 4 SAN switch connections for VSP E1090 Fibre Channel ports.
- 4 SAN switch connections for management server HBA ports.

The following table lists the Fibre Channel switch connection configuration on the VSP E1090 and database servers.

Server	HBA Ports	Storage Host Group	Switch Zone	Storage Port	Brocade G720S
BullSequana SH20 Server 1	HBA1_1	SH20_62_HBA1_1	SH20_62_HBA1_1_VSP_E 1090_33_3A	3A	SAN-switch1
	HBA1_2	SH20_62_HBA1_2	SH20_62_HBA1_2_VSP_E 1090_33_4A	4A	SAN-switch2

Server	HBA Ports	Storage Host Group	Switch Zone	Storage Port	Brocade G720S
	HBA2_1	SH20_62_HBA2_1	SH20_62_HBA2_1_VSP_E1090_33_3B	3B	SAN-switch1
	HBA2_2	SH20_62_HBA2_2	SH20_62_HBA2_2_VSP_E1090_33_4B	4B	SAN-switch2
BullSequana SH20 Server 2	HBA1_1	SH20_64_HBA1_1	SH20_64_HBA1_1_VSP_E1090_33_7A	7A	SAN-switch1
	HBA1_2	SH20_64_HBA1_2	SH20_64_HBA1_2_VSP_E1090_33_8A	8A	SAN-switch2
	HBA2_1	SH20_64_HBA2_1	SH20_64_HBA2_1_VSP_E1090_33_7B	7B	SAN-switch1
	HBA2_2	SH20_64_HBA2_2	SH20_64_HBA2_2_VSP_E1090_33_8B	8B	SAN-switch2
HA810 G3 Server 1	HBA1_1	HA810_G3_68HBA1_1	HA810_G3_68_HBA1_1_ASE32_50_1A	1A	SAN-switch1
	HBA1_2	HA810_G3_68_HBA1_2	HA810_G3_68_HBA1_2_ASE32_50_2A	2A	SAN-switch2
HA810 G3 Server 2	HBA1_1	HA810_G3_69_HBA1	HA810_G3_69_HBA1_1_ASE32_50_1B	1B	SAN-switch1
	HBA1_2	HA810_G3_69_HBA2	HA810_G3_69_HBA1_2_ASE32_50_2B	2B	SAN-switch2



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

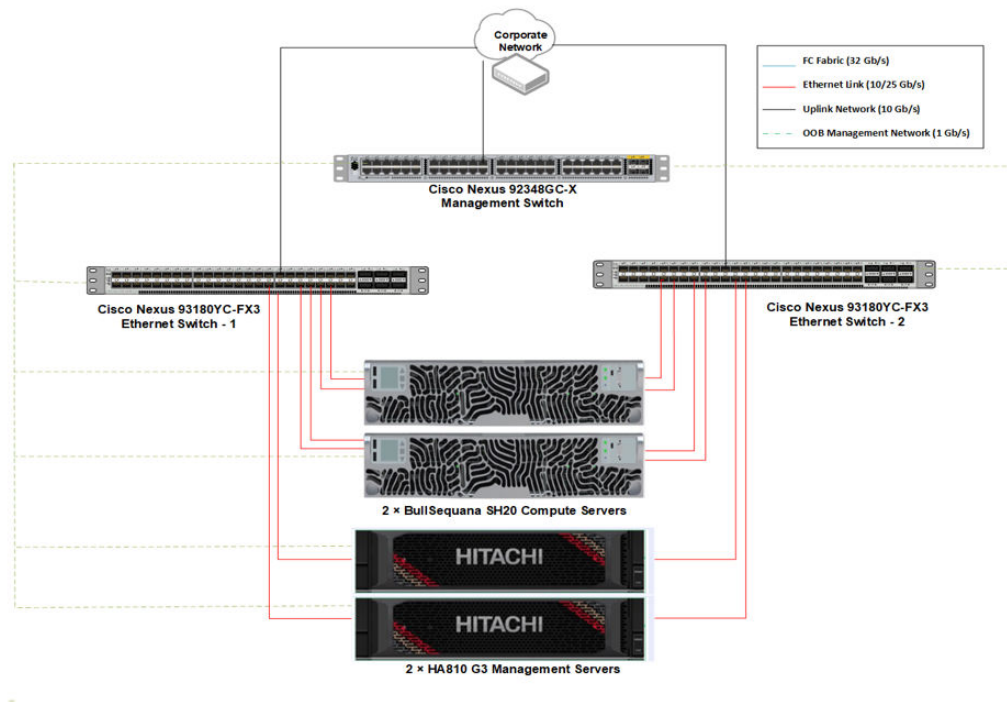
This architecture uses the following separate networks:

- 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.

Note that Hitachi Vantara recommends using pairs of 25 Gbps NICs for the cluster interconnect network and public network. We recommend following 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.

The following figure shows the IP network switch connection.



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

When creating NIC bonding pairs, paired ports should be on different cards to eliminate a single point of failure (SPoF).

Server	NIC Port	VLAN/ Subnet	NIC Bond	IP Address	Network	Bandwidth (Gbps)	Cisco Nexus 93180YC-FX3 switch	
							Switch Number	Port
Oracle Database Server1	NIC 1 PORT 1	208	Bond0	192.168.208.91	Private	25	1	31
	NIC 2 PORT 1					25	2	
	NIC 1 PORT 2	242	Bond1	192.168.242.91	Public Oracle	25	1	32
	NIC 2 PORT 2					25	2	
	iLo- Dedicated NIC	242	-	192.168.242.161	Manage- ment	1		
Oracle Database Server2	NIC 1 PORT 1	208	Bond0	192.168.208.92	Private	25	1	34
	NIC 2 PORT 1					25	2	
	NIC 1 PORT 2	242	Bond1	192.168.242.92	Public Oracle	25	1	35
	NIC 2 PORT 2					25	2	
	iLo- Dedicated NIC	242	-	192.168.242.162	Manage- ment	1		
Manage- ment Server 1	iLo- Dedicated NIC	242	-	192.168.242.169	Manage- ment	1		
	NIC 1 PORT 1	244	-	192.168.244.101	Public Oracle	25	1	

Server	NIC Port	VLAN/ Subnet	NIC Bond	IP Address	Network	Bandwidth (Gbps)	Cisco Nexus 93180YC-FX3 switch	
							Switch Number	Port
Management Server 2	iLo- 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 address and SCAN name configuration that was used when testing the environment.

Server	Virtual IP Address	Scan Name pub-scan
Database Server 1	192.168.242.91	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.

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 validation test for the Hitachi Unified Compute Platform architecture for Oracle RAC deployment with Hitachi Virtual Storage Platform E1090.

Test methodology

Oracle Orion and the peakmarks 9.6 benchmark software tools were 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.

Oracle Database Parameter	Value
Compatible	19.19.0.0.0
cluster_database	TRUE
cluster_database_instances	2
Oracle Database size	16 TB
Database storage type	ASM

Test environment

The 2 Node Oracle RAC database with BullSequana SH20 and Hitachi Virtual Storage Platform E1090 storage was configured using the architecture mentioned in this document.

We used 64 raw LUNs each of 250 GB for Orion testing. The peakmarks database was 16 TB consisting of 181 ASM disks as part of various diskgroups, and peakmarks test cases were executed for various workloads using this database.

The following table lists the test environment details.

Item	Description	Value
Operating system	-	OEL 9.1 UEK7U1
Workload type	-	OLTP/OLAP
Database size	-	16TB
Number of physical cores	CPU cores	120 per node
Memory	Server memory	2048 GB per node
Network	Cluster interconnect	2 × 25 Gbps NIC bonding

Test results

Hitachi UCP for Oracle databases solutions is built using Hitachi VSP E1090 storage delivering highly reliable, scalable and low-microsecond latency and Hitachi advance server delivering with latest generation processor, networking capabilities and large memory.

This database solution delivers unprecedented performance, scalability and availability for all possible database workloads. This unified solution with enhanced performance, highly scalable storage and memory capacity foster the opportunities to consolidate the databases with different workloads and dramatically reduce the cost of hardware and software stacks.

The new generation solution is optimized with the latest storage capabilities, network enhancement, emerging oracle database software functionalities and highly efficient compute which can deliver unmatched performance for data analytics, machine learning, AI capabilities, and online transaction systems, as well as mixed database workloads.

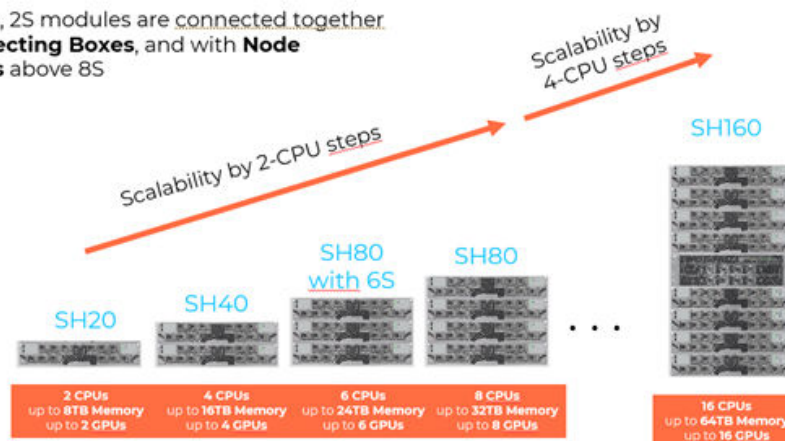
The cost-effective architecture of this solution featured with lower latency, proven availability and scalability along with unparalleled performance make this solution best for all businesses in on-premises, hybrid-cloud environments.

We have implemented this Hitachi solution for 2 node Oracle RAC database using 2 × BullSequana SH20 servers. Customers who require expanding their Oracle RAC database infrastructure beyond 2 nodes can leverage this modular and scalable architecture to scale-out their environment using Eviden BullSequana SH40/SH60/SH80 servers, respectively. See the following scale-out configuration for Eviden BullSequana SH series servers.

BullSequana SH series concept

Scalability with 2-sockets modules building blocks

To scale up, 2S modules are connected together with **Connecting Boxes**, and with **Node Controllers** above 8S



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

Category	Key Performance Metric	peakmarks® Workload
	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

Category	Key Performance Metric	peakmarks® Workload
	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

Product descriptions

The following products are used in this solution.

Hitachi Virtual Storage Platform E1090

The [Hitachi Virtual Storage Platform E1090 \(VSP E1090\)](#) storage system is a high-performance, large-capacity data storage system. The VSP E1090 all-flash arrays (AFAs) support NVMe and SAS solid-state drives (SSDs). The VSP E1090H hybrid models can be configured with both SSDs and hard disk drives (HDDs).

- The NVMe flash architecture delivers consistent, low-microsecond latency, which reduces the transaction costs of latency-critical applications and delivers predictable performance to optimize storage resources.
- The hybrid architecture allows for greater scalability and provides data-in-place migration support.

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.

Eviden BullSequana SH20

Eviden BullSequana SH20 is a high-performance two-socket rackmount server designed for optimal performance and power efficiency. This allows owners to upgrade computing performance without overextending power consumption and offers non-latency support to virtualization environments that require the maximum memory capacity. BullSequana SH series server provides flexible I/O scalability for today's diverse data center application requirements.

Optimized for performance, high density, and power efficiency in a dual-processor server, BullSequana SH20 servers deliver a balance of compute and storage capacity. These rack mounted servers have the flexibility to power a wide range of solutions and applications.

The highly scalable memory for BullSequana SH20 supports up to 4 TB RAM per module using 32 slots of 4800 MHz DDR5 RDIMM, however memory can be upgrade to more than 16 TB of RAM if upgraded with additional modules. BullSequana SH20 is powered by the Intel Xeon platinum processor family for complex and demanding workloads. Flexible OCP and PCIe I/O expansion card options are 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

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

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.

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.

Hitachi Vantara

Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA



HitachiVantara.com/contact