

Hitachi Solution for Databases – Oracle Real Application Clusters Virtualized on Oracle Linux Virtualization Manager with HA820 G2 and Virtual Storage Platform E1090

Reference Architecture Guide

© 2024 Hitachi Vantara LLC. All rights reserved.

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

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

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

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

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

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

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

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

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

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

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

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

Feedback

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

Revision history

Changes	Date
Updated infrastructure and physical illustrations.	February 16, 2024
Updated peakmarks information.	November 14, 2022
Updated <i>Business benefits</i> section.	September 8, 2022

Reference Architecture Guide

Hitachi Solution for Databases with Oracle Real Application Clusters (RAC) virtualized on Oracle Linux Virtualization Manager (OLVM) using Virtual Storage Platform E1090 and Hitachi Advanced Server HA820 G2 with Intel Xeon 3rd generation Scalable Processors provides reliability, high availability, and scalability while processing small to large Oracle workloads. This reference architecture provides a virtualized Oracle RAC Database 19c on Oracle Linux Virtualization Manager 4.3.10. In this solution, Oracle RAC Database nodes are deployed on Oracle KVM hosts with Oracle Linux 7.9 UEK 6.3 and test virtual machines (VMs) with Oracle Linux 8.5 UEK 6.3.

Hitachi Solution for Databases with Oracle RAC Virtualization Manager is a fast, agile offering from Hitachi Vantara that makes efficient use of deployed resources. The goal of Oracle RAC virtualization using Oracle Linux Virtualization Manager is IT transformation because you want a fast and flexible journey to reduce costs and quickly scale your environment up or down.

Oracle Linux Virtualization Manager release 4.3.10 supports multiple hosts running Oracle Linux KVM. The heart of the OLVM is the oVirt Engine, which is used to discover KVM hosts and configure storage and networking for the virtualized data center. This solution integrates many innovative technologies from Hitachi Vantara and Oracle. To create an Oracle RAC Virtualization Manager environment, the following components are used:

- Hitachi Virtual Storage Platform E1090 (VSP E1090) used as storage resource
- Hitachi Advanced Server HA820 G2 with Intel® Xeon® Platinum 8368 Processor 38C CPUs for storage and computing resources
- Hitachi Advanced Server DS120 G2 with Intel® Xeon® Silver 4310 Processor 12-core CPUs for the management servers
- Oracle Linux Virtualization Manager 4.3.10
- Oracle Real Application Clusters (RAC)

Running Oracle RAC on Oracle Linux Virtualization Manager the provides following benefits:

- Reduced overhead for the hypervisor to minimize infrastructure and software licensing costs
- Simplified management with Oracle Linux Virtualization Manager
- Easier to move VMs to the cloud

This solution provides the flexibility to select storage and compute resources based on unique requirements. You can 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 professional responsible for planning and deploying an Oracle Database solution

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

- Hitachi Virtual Storage Platform E1090
- Hitachi Advanced Server HA820 G2 servers
- Hitachi Advanced Server DS120 G2 servers
- Storage Area Networks
- Oracle Database administration
- Oracle Database 19c Release 15
- Oracle Linux
- Oracle Linux Virtualization Manager



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

Use this reference architecture to implement Hitachi Solution for Databases with Oracle RAC virtualized on Oracle Linux Virtualization Manager using Virtual Storage Platform E1090 and Hitachi Advanced Server HA820 G2. This solution is engineered, pre-tested, and qualified to provide high performance and high reliability in demanding, dynamic Oracle environments.

Business benefits

Oracle Linux Virtualization Manager provides the foundation for a cloud infrastructure. It provides fully integrated enterprise management from disk to cloud applications. This virtualization management platform can be easily deployed to configure, monitor, and manage an Oracle Linux Kernel-based Virtual Machine (KVM) environment with enterprise-grade performance and support from Oracle.

Here are some benefits of this reference architecture:

- This solution provides performance results close to bare metal by innovative tuning of the Oracle KVM High Performance option.
- Oracle license cost reduction with dynamic capacity planning to use a subset of CPU cores in the database server for easy configuration of Oracle KVM hard partitioning.
- Create Oracle Linux Virtualization Manager clusters with multiple Hitachi Advanced Server HA820 G2 servers to deploy virtual machines for an Oracle database deployment.
- Reduce operation and support costs while increasing IT efficiency and agility on premises and in the cloud with Oracle Virtualization.
- Add or remove servers in the Oracle Linux Virtualization Manager pool based on business requirements.
- Achieve high Oracle Database performance with VSP E1090 storage systems.
- Provide a solution for customers who are looking for converged Oracle products not covered by Hitachi VSP G700, VSP G800, VSP G900, or VSP 5000 for Oracle RAC.

High-level infrastructure

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

- Hitachi Advanced Server HA820 G2 servers
- Hitachi Advanced Server DS120 G2 servers
- Hitachi Virtual Storage Platform E1090
- Hitachi Virtual Storage Platform E990
- Brocade G720 32 Gbps SAN infrastructure
- Cisco 10/25 GbE LAN infrastructure




Note: The VSP E990 is not required in your environment. Data volumes can also be configured on the VSP E1090 or other VSP storage systems.

The Virtual Storage Platform E1090 and Hitachi Advanced Server HA820 G2 configuration has the following characteristics:

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

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 \(on page 26\)](#).

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

Vendor	Hardware	Detail Description	Version	Quantity
Hitachi Vantara	VSP E1090	6 × CHA pairs (8 × 32 Gbps Fibre Channel ports in use) 1024 GB cache memory 49 x 1.9 TB NVMe SDDs*	93-06-01-80/00	1
	VSP E990 – Management node storage system	Two Controllers 4 × 32 Gbps Fibre Channel ports 16 × 12 Gbps backend SAS ports 1024 GB cache memory 32 × 3.8 TB NVMe SSDs*	93-02-02-60/00	1
Hitachi Vantara	HA820 G2	2 × Intel Xeon Platinum 8368 38C CPUs @ 2.40GHz 768 GB (32GB × 24) DIMM DDR4-3200 Synchronous Registered (Buffered) 3200 MHz	System ROM: U46 v1.58 (01/13/2022) iLO 5: 2.65 Feb 23 2022	2
		2 × HPE Ethernet 10/25 Gb 2-port FLR-SFP28 BCM57414 adapters	Driver: bnxt_en Driver version: 1.10.2 (inbox) Firmware: 219.0.144.0	
		2 × HPE SN1610E 32 Gb 2 port Fibre Channel HBAs	Driver: lpfc Driver version: 12.8.0.11 Firmware: 12.8.528.21	
	DS120 G2	2 × Intel Xeon Processors 4310, 12-core, 2.10 GHz, 120W 256 GB (32 GB × 8) DIMM DDR4-3200 Synchronous Registered (Buffered) 3200 MHz	BIOS: S5XH3A14.H00 BMC: 3.32.06 CPLD: 07	2

Vendor	Hardware	Detail Description	Version	Quantity
		1 × 256 GB NVMe 0.3DWPD M.2 SSD for boot		
		1 × Dual Port 25 GbE NIC Intel E810 PCIe card	Driver version: ice 0.8.2-k Firmware: 2.42	
		1 × Emulex LightPulse LPe35002-M2 2-Port 32 Gb Fibre Channel Adapter	Driver: lpfc Driver Version: 12.8.0.10 Firmware: 12.8.542.26	
Brocade	G720 Fibre Channel switch	48 × 32 Gbps port Fibre Channel switches 32 Gbps SFPs	Fabric OS: v9.0.1c	2
Cisco	Cisco Nexus N9K- C93180YC- FX3	48 × 10/25 GbE ports 6 × 40/100 Gbps Quad SFP (QSFP28) ports	BIOS: version 05.23 NXOS: version 7.0(3)I7(1)	2
	Cisco- C92348GC-X	48 × 1 GbE port Ethernet switches	BIOS: version 05.40 NXOS: version 9.3(2)	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.



Note: You can choose larger capacity SSDs to fit their business requirements.

The following table lists the software components for the compute nodes.

Software	Version	Function
Oracle Linux	7.9 with 5.4.17-2102.203.6.el7uek.x86_64 on OLVM 7.9 with 5.4.17-2136.307.3.6.el7uek.x86_64 on KVM hosts 8.5 with 5.4.17-2136.306.1.3.el8uek.x86_64 on RAC VMs	Operating system
Oracle Database	19c (Version 19.15.0.0.0)	Database software
Oracle Real Application Clusters	19c (Version 19.15.0.0.0)	Cluster software
Oracle Grid Infrastructure	19c (Version 19.15.0.0.0)	Volume management, file system software, and Oracle automatic storage management

The following table lists the software components for the management nodes.

Software	Version	Function
ESXi	7.0 U2	Hypervisor for management server
vCenter Server	7.0 U2	VMware cluster management server
Hitachi Device Manager - Storage Navigator	Microcode-dependent	Storage management software
Hitachi Ops Center	10.8.3	Hitachi infrastructure management software
Oracle Linux Virtualization Manager	4.3.10	Server virtualization management platform

Solution design

This section describes the reference architecture environment to implement Hitachi Solution for Databases with Oracle RAC virtualized on Oracle Linux Virtualization Manager. The environment uses Hitachi Virtual Storage Platform E1090.

The infrastructure configuration includes the following:

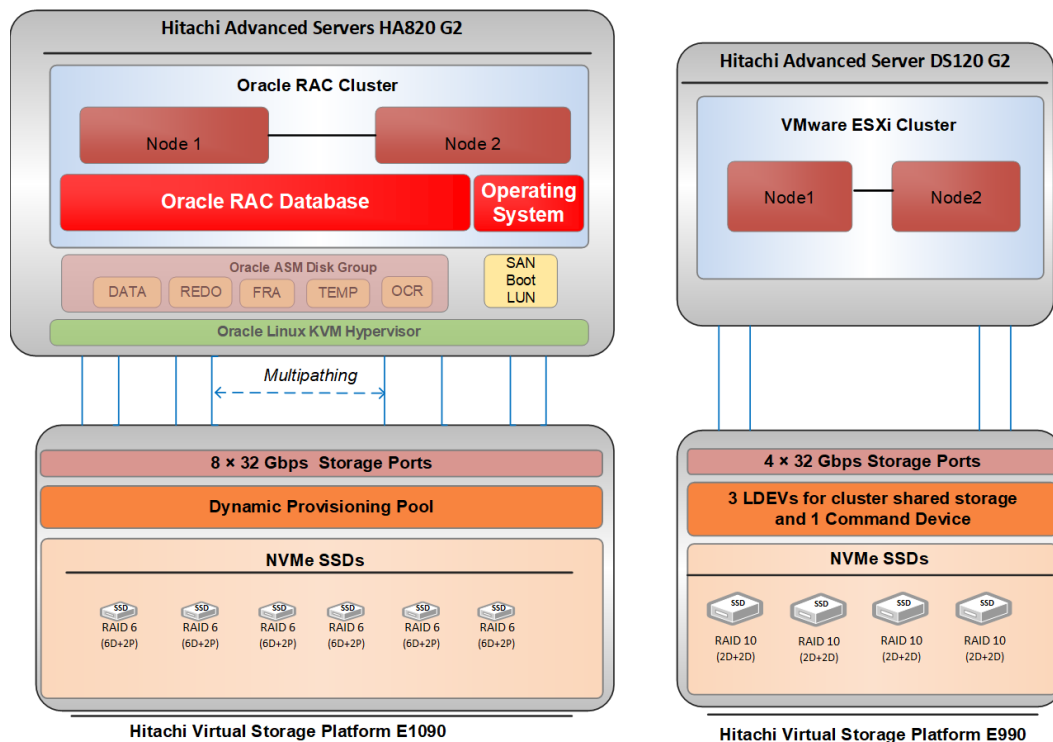
- Oracle KVM hosts — There are two hosts configured in an Oracle database environment.
- Management Node — One Oracle Linux Virtualization Manager node is used to manage KVM hosts.
- Storage System — There are vVols mapped to each port that are presented to the server as LUNs.
- SAN Connection — There are SAN connections to connect the Fibre Channel HBA ports to the storage through Brocade G720 switches.

Storage architecture

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

Storage configuration

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



The following table shows the VSP E1090 storage pool configuration used for this solution.

Pool ID	Oracle RAC-Pool
Pool Type	Dynamic Provisioning Pool
RAID Group	1-1 to 1-6
RAID Level	RAID 6 (6D+2P)
Drive Type	1.9 TB SSD
Number of Drives	48
Number of LDEVs	48
LDEV Sizes	1320 GB
Pool Capacity	61.89 TB

In this solution, servers use SAN boot with RAID 6 protection. The following table shows the VSP E1090 logical storage configuration.

ASM disk group	Total number of dynamic provisioning volumes	Dynamic provisioning volume sizes (GB)	Purpose	Storage Ports
DATA	16	200 GB	OLTP Application Tablespaces, System Sysaux, Undo	1A, 2A, 3A, 4A, 5A, 6A, 7A, 8A
REDO	8	20 GB	Online Redo Logs and Control Files	
TEMP	4	200 GB	Temp Tablespace	
FRA	4	200 GB	Incremental Backups, Archived Redo Logs, Control File and Auto backups	
OCR	3	15 GB	Oracle Cluster Registry and Voting Disk	
N/A	2	380 GB	SAN Boot OS Volumes	

The following table lists the details of the VSP E990 configuration for management servers.

Item	Value/Description
Purpose	VMware datastores CCI device
RAID Level	RAID 6 (6D+2P)
Drive Type	1.9 TB SSD
Number of Drives	16
Number of Spare Drives	0
Number of LDEVs	10
LDEV Sizes	2048 GB
Number and Size of CCI device	1 × 100 MB
Storage Port for Management Servers	5C, 6C

Additional LUNs can be mapped if required. While the test environment was configured using a dedicated SAS RAID group for the management server cluster, it can be configured as a dedicated SSD RAID group, a dedicated dynamic provisioning pool, or as 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 E1090 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 the 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 diskgroup with external redundancy for the data and indexes tablespace.
- TEMP Tablespace — Place the TEMP tablespace in this configuration in the TEMP ASM diskgroup.
- Undo Tablespace — Create an UNDO tablespace in this configuration within the Oracle Data ASM diskgroup.
- Online Redo Logs — Create an ASM diskgroup with external redundancy for Oracle online redo logs.
- Oracle Cluster Registry and Voting Disk — Create an ASM diskgroup with normal redundancy to contain the OCR and voting disks and to protect against single disk failure to avoid loss of cluster availability. Place each of these files in this configuration in the OCR ASM diskgroups.

Oracle initial parameter settings

The following table shows the Oracle Database settings.

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

Oracle ASM disk mapping

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 the parameters accordingly when multiple virtual machine pairs are used.

ASM Disk Group	ASM Disk	UDEV Rules	LUN Details	Purpose
N/A	N/A	/dev/xvd[a-b]1	2 × 380 GB	OS and Oracle Database
OCRDG	OCR1-OCR3	/dev/xvd[c-e]1	3 × 15 GB	Oracle cluster registry and voting disk
DATADG	DATA1-DATA16	/dev/xvd[f-u]1	16 × 200 GB	Application data
REDODG	REDO1-REDO8	/dev/xvd[v-z]1, /dev/xvda[a-c]1	8 × 20 GB	Online REDO log group
FRADG	FRA1-FRA4	/dev/xvda[d-e]1	4 × 200 GB	Flash recovery area
TEMPDG	TEMP1-TEMP4	/dev/xvda[f-i]1	4 × 200 GB	Temporary tablespace

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 used two Hitachi Advanced Server HA820 G2 servers with 3rd Generation Intel Xeon Scalable Processors for each VSP storage system architecture that was tested.

These servers provide the compute power for the Oracle database to manage complex database queries and to process large volumes of transaction processing in parallel.

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

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
HA820 G2 KVM Host	KVM server-01	Oracle Linux KVM Hypervisor Host (Version 4.2.1) (KVM host1)	2 × Intel Xeon Platinum 8368 38C CPUs	152 (2 × 2 × 38)	768 GB (32 GB × 24)
	KVM server-02	Oracle Linux KVM Hypervisor Host (Version 4.2.1) (KVM host2)		152 (2 × 2 × 38)	768 GB (32 GB × 24)
KVM VMs	VM1	Oracle Linux KVM VM for Oracle RAC Node 1	2 × Intel Xeon Platinum 8368 38C CPUs	2 × 32	128 GB
	VM2	Oracle Linux KVM VM for Oracle RAC Node 2			
DS120 G2	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 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 HPE SN1610E HBAs per Advanced Server HA820 G2.

Compute servers

- 8 SAN switch connections are used for VSP E1090 Fibre Channel ports.
- 8 SAN switch connections are used for server HBA ports.

Management servers

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

The following table lists the details of the Fibre Channel switch connect configuration on Hitachi VSP E1090 ports.

Server	HBA	Host Group Name	Host Name	Switch Zone	Storage System	Storage Port	Brocade G720 Switch
HA820 G2 Server 1	HBA1	CN31	CN31_HBA1_1	CN31_HBA1_1_ASE42_43_1A	VSP E1090	1A	31
	HBA2	CN31	CN31_HBA1_2	CN31_HBA1_2_ASE42_43_2A		2A	32
	HBA3	CN31	CN31_HBA2_1	CN31_HBA2_1_ASE42_43_3A		3A	31
	HBA4	CN31	CN31_HBA2_2	CN31_HBA2_2_ASE42_43_4A		4A	32
HA820 G2 Server 2	HBA1	CN32	CN32_HBA1_1	CN32_HBA1_1_ASE42_43_5A		5A	31
	HBA2	CN32	CN32_HBA1_2	CN32_HBA1_2_ASE42_43_6A		6A	32
	HBA3	CN32	CN32_HBA2_1	CN32_HBA2_1_ASE42_43_7A		7A	31
	HBA4	CN32	CN32_HBA2_2	CN32_HBA2_2_ASE42_43_8A		8A	32

The following table lists the details of the Fibre Channel switch connect configuration on Hitachi VSP E990 ports.

Server	HBA	Host Group Name	Host Name	Switch Zone	Storage Port	Brocade G720 Switch
DS120 G2 Server 1	HBA1	MN33	MN33_HBA1	MN33_HBA1_ASE42_43_5C	5C	69
	HBA2	MN33	MN33_HBA2	MN33_HBA2_ASE42_43_6C	6C	70
DS120 G2 Server 2	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, it is recommended to 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

Do the following when configuring networks in your environment:

- Use NIC bonding to provide failover and load balancing of interconnections within a server.
- Set all NICs to full duplex mode.

Configure each Oracle KVM server node with at least the bonding interfaces for the following:

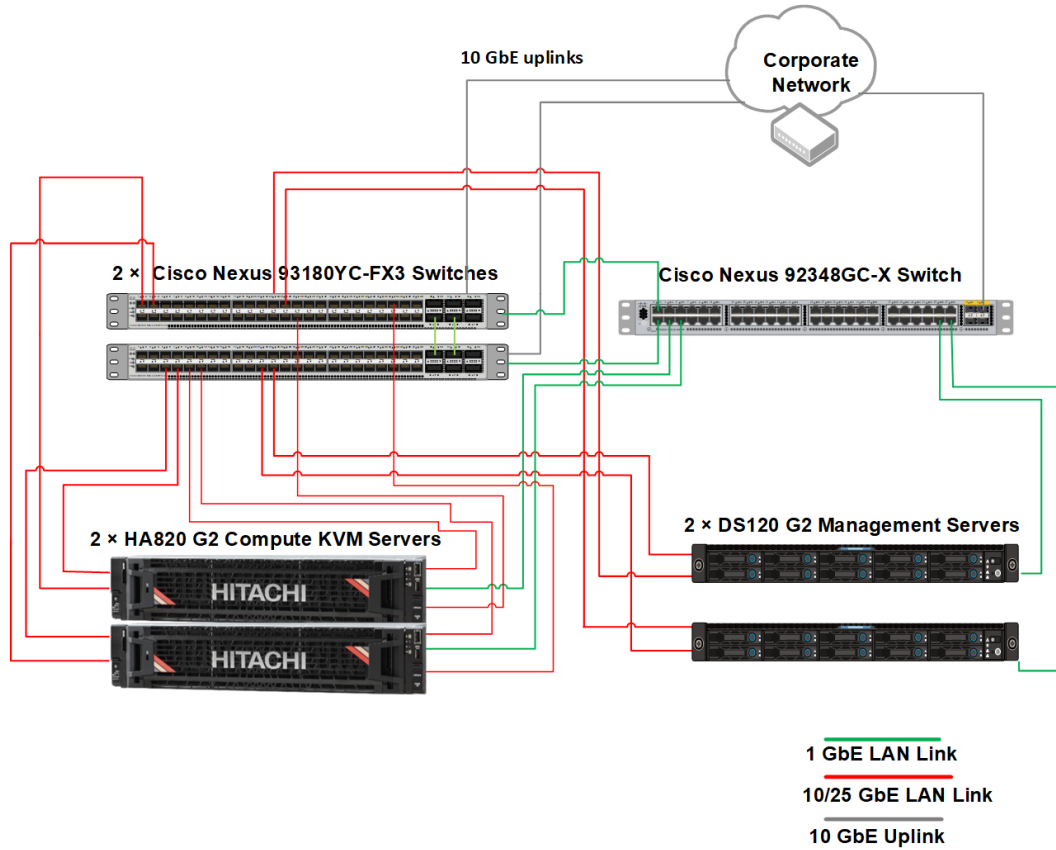
- Management network
- Public network
- Private network

Each virtual machine has public, private, and management vNICs. Use separate VLANs for the following:

- Oracle KVM management network
- Private Oracle Database network
- Public network

Physical network configuration

The following figure shows the physical IP network switch connections.



The following table lists the HA820 G2 and DS120 G2 network configuration for this solution.

Server	NIC Ports	Subnet	NIC Bond	IP Address	Network	Bandwidth (Gb/s)	Cisco Nexus 93180YC-FX 3 Switch	
							Switch Number	Port
HA820 G2 Server 1	NIC 1-Port 0	242	Bond0	192.198.242.204	Management/Public	10	1	41
	NIC 2-Port 1					10	2	
	NIC 2-Port 0	100	Bond1	192.168.100.xx	Private	25	1	42
	NIC 1-Port 1					25	2	
HA820 G2 Server 2	NIC 1-Port 0	242	Bond0	192.198.242.206	Management/Public	10	1	43
	NIC 2-Port 1					10	2	
	NIC 2-Port 0	100	Bond2	192.168.100.xx	Private	25	1	44
	NIC 1-Port 1					25	2	
DS120 G2 Management Server 1	NIC 1-Port 1	242	Bond0	192.198.242.101	Management/Public	25	1	49

Server	NIC Ports	Subnet	NIC Bond	IP Address	Network	Bandwidth (Gb/s)	Cisco Nexus 93180YC-FX 3 Switch	
							Switch Number	Port
	Port 0							
	NIC 1-Port 1					25	2	
DS120 G2 Management Server 2	NIC 1-Port 0	242	Bond0	192.198.242.102	Management/Public	25	1	50
	NIC 1-Port 1					25	2	



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 the servers and VSP E1090.

Name	IP Address
Oracle Linux KVM host 1	192.168.242.xx
Oracle Linux KVM host 2	192.168.242.xx
Management Server 1	192.168.242.xx
Management Server 2	192.168.242.xx
VSP E1090	192.168.242.xx
VSP E1090 CTL1	192.168.242.x
VSP E1090 CTL2	192.168.242.x

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 this solution.

Database configuration

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

Oracle Database Parameter	Value
Compatible	19.15.0.0.0
Oracle Database size	2 TB
Database storage type	ASM
Database fill factor	80%

Test environment

The following table lists the configuration details for VSP E1090 testing.

Item	Value
Operating System on VMs	OL 8.5
Workload Type	OLTP/OLAP
Database Size	2 TB
Number of vCPUs	32
Virtual Memory	128 GB per VM
Host Cluster VM Network	2 × 10 Gbps NIC Bonding

Test methodology

Oracle Orion and peakmarks OLTP test software was used to validate this solution.

Orion

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

For more information about Orion, see “I/O Configuration and Design” in the Oracle Database Performance Tuning Guide at <https://docs.oracle.com/en/database/oracle/oracle-database/12.2/tgdba/IO-configuration-and-design.html#GUID-B89C41F4-B12A-49E1-BFB9-5BF978E2D6B3>.

peakmarks

Test results are demonstrated using peakmarks OLTP test cases.

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

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

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

Test results

The following table lists Orion test result details.

Test Category	Test Case	Test Result	
		Throughput	RT (ms)
Orion	8 KB Random Read	395,332 IOPS	1.09
	8 KB Random Write	169,052 IOPS	0.83
	1 MB Sequential Read	21.89 GBps	N/A
	1 MB Sequential Write	9.31 GBps	N/A

The following table provides the peakmarks test results. The peakmarks 9.6 tool was used to validate this solution.

There are two types of VMs that were configured for Oracle Real Application Clusters Virtualized on Oracle Linux Virtualization Manager, namely VMs optimized for "Server Class VM" and VMs optimized for "High Performance VM."

We compared the performance results in the next section for different Oracle Database workloads with the standard recommended configuration along with CPU hard partitioning for Server Class VM and High Performance VM. The test results were for random, transactional, analytics, and database background-related process workloads with High Performance VM which is closer to bare metal performance.

Therefore, we recommend using High Performance VM over Server Class VM. For recommended best practices and configuration see https://access.redhat.com/documentation/en-us/red_hat_virtualization/4.2/html/virtual_machine_management_guide/configuring_high_performance_virtual_machines_templates_and_pools#Automatic_High_Performance_Configuration_Settings.

peakmarks Workload	Test Results			
	Server Class VM		High Performance VM	
	Throughput	RT (ms)	Throughput	RT (ms)
Storage Performance				
STO-READ SQL sequential read	23.24 GBps	N/A	23.70 GBps	N/A
STO-RANDOMSQL statements random reading/updating blocks via index access, 100% read	390,359 IOPS	0.86	879,687 IOPS	0.86
STO-RANDOM SQL statements random reading/updating blocks via index access, 80% read	362,843 IOPS	0.736	718,668 IOPS	0.78
STO-RANDOMSQL statements random reading/updating blocks via index access, 50% read	327,673 IOPS	0.99	452,323 IOPS	0.51
Transaction Processing Performance				
TP-LIGHT light transaction, select/update single row via index per transaction 1 rpt 0%update	483,365 tps	0.23	921,845 tps	0.21
TP-LIGHT light transaction, select/update single row via index per transaction 1rpt 20 %update	250,189 tps	0.4	280,041 tps	0.53
TP-LIGHT light transaction, select/update single row via index per transaction 1rpt 30%update	203,605 tps	0.49	237,590 tps	0.54
TP-MEDIUM transaction, select/update 25 rows via index per transaction 25 rpt 0%update	24,797 tps	0.97	74,935 tps	0.93
TP-MEDIUM transaction, select/update 25 rows via index per transaction rpt 20%update	14,731 tps	1.62	17,680 tps	0.90
TP-MEDIUM transaction, select/update 25 rows via index per transaction 25 rpt 30%update	12,591 tps	1.27	12,702 tps	1.26

peakmarks Workload	Test Results			
	Server Class VM		High Performance VM	
	Throughput	RT (ms)	Throughput	RT (ms)
TP-HEAVY transaction, select/ update 125 rows via index per transaction 125 rpt 0%update	N/A	N/A	35,108 tps	0.51
TP-HEAVY transaction, select/ update 125 rows via index per transaction 125 rpt 20%update	N/A	N/A	9,678 tps	1.65
TP-HEAVY transaction, select/ update 125 rows via index per transaction 125 rpt 30%update	N/A	N/A	8,291 tps	1.92
Data Analytics Performance				
DA-STORAGE data scan storage system	23.79 GBps	N/A	23.96 GBps	N/A
DA-ROWSTORE data scan row store	104.10 GBps	N/A	104.42 GBps	N/A
Database background processes Performance				
DBWR-THR throughput of changed blocks written back to storage system	2.12 GBps	N/A	2.48 GBps	N/A
LGWR-LAT1 commit rate and commit latency for small size transaction, 1 rpt	N/A	N/A	131,944 tps	0.83
LGWR-THR throughput of redo data written by log writer processes	0.92 GBps	N/A	1.63 GBps	N/A
Server Performance				
SRV-SCAN	149.65 GBps	N/A	143.54 GBps	N/A
SRV-QUERY1 throughput for simple queries, select 1 row via index	N/A	N/A	6,139,889 qps	0.013
SRV-QUERY25 throughput for complex queries, select 25 rows via index	N/A	N/A	1,408,334 qps	0.057

peakmarks Workload	Test Results			
	Server Class VM		High Performance VM	
	Throughput	RT (ms)	Throughput	RT (ms)
SRV-REPORT logical read throughput, select 125 rows via index	N/A	N/A	43,439,986 dbps	N/A
SRV-MIXED throughput for mixed queries and full table scans	N/A	N/A	2,023,957 qps	0.032
Data load performance				
DL-BUFFER transactional data load via buffer cache	N/A	N/A	498 MBps	N/A

Conclusion

Hitachi Virtual Storage Platform E1090 and Hitachi Advanced Server HA820 G2 were tested and validated as an ideal platform for virtualized workloads because the solution delivers efficient performance, cost effectiveness, rapid provisioning, and sub-capacity licensing. You can configure and use multiple Oracle clusterwares, other application clusters, or individual VMs that can be used for various environments over KVM hosts, along with Oracle virtualized clusters, thus efficiently using all system resources.

Product descriptions

These products are used in this reference architecture.

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 Advanced Server HA820 G2

Hitachi Advanced Server HA820 G2 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 maximum memory capacity. Hitachi Advanced Server HA820 provides flexible I/O scalability for today's diverse data center application requirements.

Hitachi Advanced Server DS120 G2

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

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

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

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 Virtualization Manager

[Oracle Linux Virtualization Manager](#) is a new server virtualization management platform that can be easily deployed to configure, monitor, and manage an Oracle Linux Kernel-based Virtual Machine (KVM) environment with enterprise-grade performance and support from Oracle.

Oracle Linux

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

Oracle Database with Real Application Clusters Option

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

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

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

VMware ESXi

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

VMware vCenter Server Appliance

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

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

Brocade switches from Broadcom

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

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

Cisco Nexus switches

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

peakmarks® test descriptions

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

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

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

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

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

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

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

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

The following is a list of measurement units:

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