

Hitachi Solution for Databases - Oracle RAC Virtualized on Oracle Linux Virtualization Manager using Virtual Storage Platform E590 and Hitachi Advanced Server DS220

Reference Architecture

Legal Notices

© 2021 Hitachi Vantara LLC. All rights reserved.

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

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

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

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

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

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

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

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

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

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

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

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

Feedback

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

Revision history

Changes	Date
Initial release	April 5, 2021

Reference Architecture

Hitachi Solution for Databases with Oracle RAC virtualized on Oracle Linux Virtualization Manager using Virtual Storage Platform E590 and Hitachi Advanced Server DS220 provides reliability, high availability, and scalability while processing small to large Oracle workloads. This reference architecture provides a virtualized Oracle RAC Database 19c on the Oracle Linux Virtualization Manager 4.3.6 solution. In this solution Oracle RAC Database nodes are deployed on Oracle KVM hosts with Oracle Linux 8.2.

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

Oracle Linux Virtualization Manager 4.3.6 supports multiple hosts running Oracle Linux KVM. The heart of the manager 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 E590 (VSP E590) used as storage resource.
- Hitachi Advanced Server DS220 with Intel Xeon Platinum 8276 processors for storage and computing resources.
- Hitachi Advanced Server DS120 with Intel Xeon Gold Processor 6240, 18-core processors for the management servers.
- Oracle Linux Virtualization Manager 4.3.6.
- Oracle Real Application Clusters (RAC).

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

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

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 E590
- Hitachi Advanced Server DS220 servers
- Hitachi Advanced Server DS120 servers
- Storage Area Networks
- Oracle Database administration
- Oracle Database 19c Release 3
- Oracle Linux
- Oracle Linux Virtualization Manager



Note: Testing of this configuration was in a lab environment. Many things affect production environments beyond prediction or duplication in a lab environment. Follow the recommended practice of conducting proof-of-concept testing for acceptable results in a non-production, isolated test environment that otherwise matches your production environment before your production implementation of this solution.

Solution overview

Use this reference architecture to implement Hitachi Solution for Databases with Oracle RAC virtualized on Oracle Linux Virtualization Manager using Virtual Storage Platform E590 and Hitachi Advanced Server DS220. This solution is engineered, pre-tested, and qualified to provide high performance and high reliability in demanding, dynamic Oracle environments. This reference architecture implements Hitachi Solution for Databases for an Oracle Real Application Clusters database using Hitachi Virtual Storage Platform E590.

Business benefits

Oracle Linux Virtualization Manager provides the foundation for the 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:

- Create Oracle Linux Virtualization Manager clusters with multiple Hitachi Advanced Server DS220 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.
- Add or remove servers in the Oracle Linux Virtualization Manager pool based on business requirements.

- Achieve high Oracle Database performance with VSP E590 storage.
- 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 DS220 servers
- Hitachi Advanced Server DS120 servers
- Hitachi Virtual Storage Platform E590
- Hitachi Virtual Storage Platform E990
- Brocade G620 32 Gbps SAN infrastructure
- Cisco 10/25 GbE LAN infrastructure

The configuration of Virtual Storage Platform E590 and Hitachi Advanced Server DS220 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.



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

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 your environment and requirements. The Cisco Nexus 93180YC-EX switches used in this reference architecture can support uplink speeds of 40 GbE or 100 GbE if higher bandwidth is required. The VSP E990 is not required in your environment. Data volumes can also be configured on the VSP E790 or other VSP storage arrays.

Key solution components

The key solution components for this solution are listed in the following tables. Detailed component information is provided in [Product descriptions \(on page 23\)](#)

Hardware components

Vendor	Hardware	Detail Description	Version	Quantity
Hitachi Vantara	Hitachi Virtual Storage Platform VSP E590	6 × CHA pairs (6 × 32 Gbps Fibre Channel ports) 624.50 GB cache memory 24 × 1.9 TB Backend NVMe SDDs	93-03-21-40/00	1
	Hitachi Virtual Storage Platform VSP E990 – Management node storage system	8 × CHA pairs (8 × 32 Gbps Fibre Channel ports) 1024 GB cache memory 16 × 1.9 TB backend NVMe SDDs	93-02-02-60/00	1
Hitachi Vantara	Hitachi Advanced DS220 Servers	2 × Intel® Xeon® Platinum 8276L CPUs @ 2.2 GHz 768 GB (64 GB × 12) DIMM DDR4 synchronous registered (buffered) 2666 MHz	BIOS: 3B18.H00 BMC: 4.67.06 CPLD: 11	2
		2 × Intel XXV710 Dual Port 25 GbE SFP28 NIC cards	Driver: i40e Driver Version: 2.8.20-k Firmware: 7.30	
		2 × Emulex LightPulse LPe32002-M2 2-Port 32 Gb Fibre Channel adapters	Driver: lpfc Driver Version: 12.6.0.3 Firmware: 12.6.240.22	
	Hitachi Advanced DS120 Servers	2 × Intel Xeon Gold Processor 6240, 18-core, 2.60GGHz, 85W 256 GB (32 GB × 8) DIMM DDR4 synchronous registered (buffered) 2666 MHz 1 × 64 GB MLC SATADOM for boot	BIOS: 3B18.H00 BMC: 4.67.06 CPLD:11	2

Vendor	Hardware	Detail Description	Version	Quantity
		1 × Dual Port 25 GbE NIC Intel XXV710 PCIe card	Driver: i40en Driver Version: 1.7.17 Firmware: 6.80	
		1 × Emulex LightPulse LPe31002-M6 2-Port 16 Gb Fibre Channel adapter	Driver: lpfc Firmware: 12.0.193.13 Driver version: 12.0.193.14	
Brocade	G620 Fibre Channel switches	24 × 48 port Fibre Channel switches 32 Gbps SFPs	Kernel: 4.1.35rt41 Fabric OS: v9.0.0b	2
Cisco	Cisco Nexus N9K-C93180YC- EX	48 × 10/25 GbE fiber ports 6 × 40/100 Gbps Quad SFP (QSFP28) ports	BIOS: version 07.65 NXOS: version 7.0(3)I7(6)	2
	Cisco- C92348CG-X	1 GE 48-Port Gb Ethernet switch	BIOS: version 5.37 NXOS: version 9.3(3)	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.

Software components for compute nodes

Software	Version	Function
Oracle Linux	7.9 with UEK 5.4.17-2011.6.2.el7uek.x86_64 for OLVM 7.9 with 5.4.17-2036.102.0.2.el7uek.x86_64 on KVM hosts	Operating system
Oracle Database	19c (Version 19.3.0.0.0)	Database software
Oracle Grid Infrastructure	19c (Version 19.3.0.0.0)	Volume management, file system software, and Oracle automatic storage management

Software components for management nodes

Software	Version	Function
ESXi	6.7.0 U3 Build 15160138	Hypervisor for management server
vCenter Server	6.7.0 U3 Build 15976714	VMware cluster management server
Hitachi Device Manager- Storage Navigator Hitachi Storage Advisor Embedded	Microcode dependent	Storage management software
Hitachi Ops Center	10.5.1	Hitachi infrastructure management software
Hitachi Adapters for Oracle Database	TBD	TBD
Oracle Linux Virtualization Manager	4.3.6	Server virtualization management platform

Solution design

This 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 E590.

The infrastructure configuration includes the following:

- Oracle KVM hosts — Two hosts are configured in an Oracle database environment.
- Management Node — One Oracle Linux Virtualization Manager node is used to manage KVM hosts.
- Storage System — vVols are mapped to each port and presented to the server as LUNs.
- SAN Connection — SAN connections from the Fibre Channel HBA ports connect the storage systems through Brocade G620 switches.

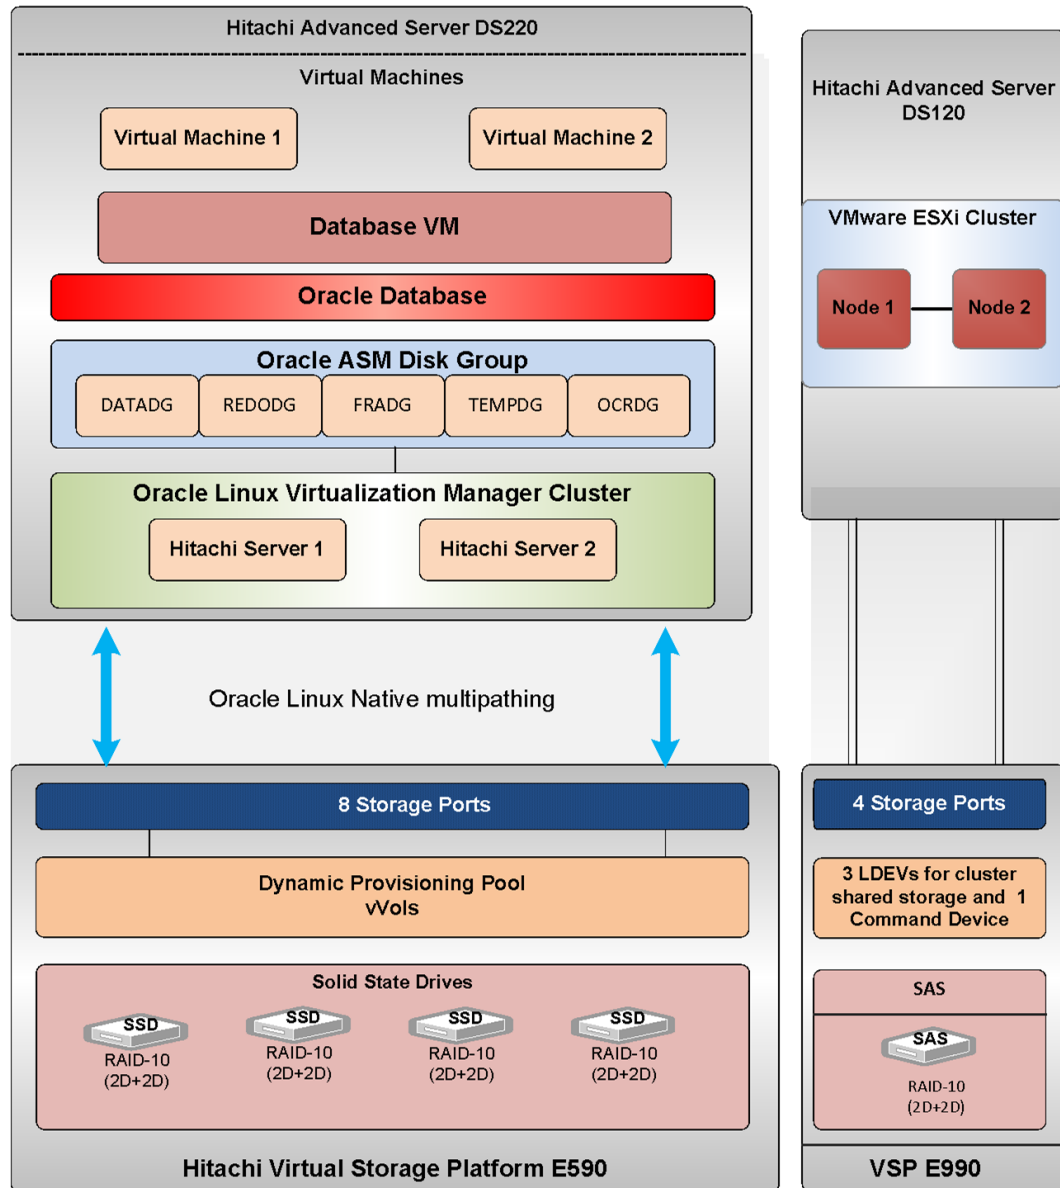
Storage architecture

This section describes the storage architecture for this solution.

Storage configuration

The storage configuration takes into consideration Hitachi 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.



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

The following table shows the storage pool configuration on the VSP E590.

Pool ID	Oracle RAC-Pool
Pool Type	Dynamic Provisioning Pool
RAID Group	1-1 to 1-6
RAID Level	RAID 10 (2D+2D)
Drive Type	1.9 TB SSD
Number of Drives	24
Number of LDEVs	24
LDEV Sizes	880 GB
Pool Capacity	20.62 TB

In this solution, servers use SAN boot with RAID 10 protection. The following table shows the VSP E590 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	16	200 GB	OLTP Application Tablespaces, System Sysaux, Undo, Temp	1A, 2A, 3A, 4A, 5A, 6A, 7A, 8A
REDO	8	10 GB	Online Redo Logs and Control Files	
TEMP	4	40 GB	Temp logs	
FRA	2	200 GB	Incremental Backups, Archived Redo Logs, Control File and Auto backups	
OCR	3	60 GB	Oracle Cluster Registry and Voting Disk	
N/A	2	200 GB	SAN Boot OS Volumes	

The following table shows 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 Ports 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, 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 E590 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 for use in various deployment scenarios to provide the right balance between performance and ease of management for a given scenario.

Oracle ASM configurations

Configure the following for Oracle ASM:

- Data and Index Tablespace — Assign an ASM diskgroup with external redundancy for the data and index tablespaces.
- TEMP Tablespace — Place the TEMP tablespace in this configuration in the TEMP ASM 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 parameters

The following table lists 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
N/A	N/A	/dev/xvd[a-b]1	2 × 200 GB	OS and Oracle Database
OCRDG	OCR1-OCR3	/dev/xvd[c-e]1	3 × 60 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 × 10 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 × 40 GB	Temporary Tablespace

Management server configurations

The following table lists management server virtual machine (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 DS220 servers with 2nd Generation Intel Xeon Scalable Processors for each storage system architecture that was tested.

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

Two Hitachi Advanced Server DS120 servers are used for 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
DS220	server-01	Oracle Linux KVM Host for Oracle RAC VM (KVM host1)	2 × Intel Xeon Platinum 8276L 28C CPUs	56 (2 × 28 C)	768 GB (64 GB × 12)
	server-02	Oracle Linux KVM Host for Oracle RAC VM (KVM host2)		56 (2 × 28 C)	768 GB (64 GB × 12)
DS120	ESXi_101	Management server for Hitachi management applications	2 × Intel Xeon Gold Processors 6240, 18-core	16	256 GB (32 GB × 8)
	ESXi_103			16	256 GB (32 GB × 8)

SAN architecture

Map the provisioned LDEVs to multiple ports on Hitachi Virtual Storage Platform E590. 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 per Advanced Server DS220.

Compute servers:

- 8 SAN switch connections are used for VSP E590 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 details of the Fibre Channel switch connection configuration on Hitachi VSP E590 ports.

Server	HBA	Host Group Name	Host Name	Switch Zone	Storage Port	Brocade G620 switch
DS220 Server 1	HBA1	CN31	CN31_HBA1_1	CN31_HBA1_1_ASE42_43_1A	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_1B	3A	31
	HBA4	CN31	CN31_HBA2_2	CN31_HBA2_2_ASE42_43_2B	4A	32
DS220 Server 2	HBA1	CN32	CN32_HBA1_1	CN32_HBA1_1_ASE42_43_3A	5A	31
	HBA2	CN32	CN32_HBA1_2	CN32_HBA1_2_ASE42_43_4A	6A	32
	HBA3	CN32	CN32_HBA2_1	CN32_HBA2_1_ASE42_43_3B	7A	31
	HBA4	CN32	CN32_HBA2_2	CN32_HBA2_2_ASE42_43_4B	8A	32

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

Server	HBA	Host Group Name	Host Name	Switch Zone	Storage Port	Brocade G620 switch
DS120 Server1	HBA1	MN33	MN33_HBA1	MN33_HBA1_ASE42_43_1A	5C	69
	HBA2	MN33	MN33_HBA2	MN33_HBA2_ASE42_43_3A	6C	70
DS120 Server2	HBA1	MN34	MN34_HBA1	MN34_HBA1_ASE42_43_2A	5C	69
	HBA2	MN34	MN34_HBA2	MN34_HBA2_ASE42_43_4A	6C	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:

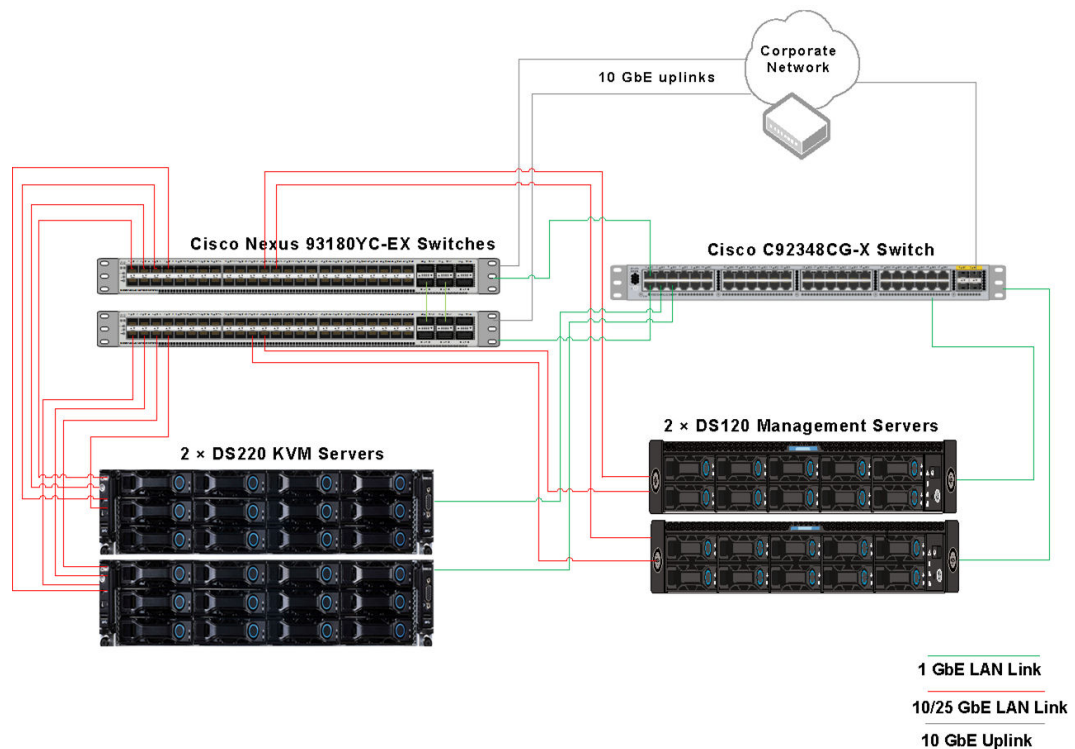
- Management network
- Public network
- Private network

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 IP network switch connections.



The following table lists the DS220 and DS120 network configuration.

Server	NIC Ports	Subnet	NIC Bond	IP Address	Network	Bandwidth (Gbps)	Cisco Nexus 93180YC-EX switch	
							Switch Number	Port
DS220 Server1	NIC 1-Port 0	242	Bond0	192.198.242.204	Management/Public	10	1	41
	NIC 1-Port 1					10	2	
	NIC 2-Port 0	242	Bond1	192.198.242.205	Public	25	1	42
	NIC 3-Port 0					25	2	
	NIC 2-Port 1	100	Bond2	192.168.100.10	Private/Public	25	1	
	NIC 3-Port 1					25	2	
DS220 Server2	NIC 1-Port 0	242	Bond0	192.198.242.206	Management/Public	10	1	43
	NIC 1-Port 1					10	2	
	NIC 2-	242	Bond1	192.198.242.207	Public	25	1	44

Server	NIC Ports	Subnet	NIC Bond	IP Address	Network	Bandwidth (Gbps)	Cisco Nexus 93180YC-EX switch	
							Switch Number	Port
	Port 0					25	2	
	NIC 3-Port 0							
	NIC 2-Port 1	100	Bond2	192.168.100.11	Private/Public	25	1	
	NIC 3-Port 1					25	2	
DS120 Management Server1	NIC 1-Port 0	242	Bond0	192.198.242.101	Management/Public	25	1	49
	NIC 1-Port 1					25	2	
DS120 Management Server2	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 servers and VSP E590.

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 E590	192.168.242.xx
VSP E590 CTL1	192.168.242.x
VSP E590 CTL2	192.168.242.x

The following table lists the network configuration for switches.

Switch Type	Model	Switch Name	IP Address for MGMT port
Cisco 1GbE Management Network Switch	Cisco Nexus C92348CG-X	C92348CG-X -1	192.168.242.xx
Cisco 10G/25GbE Network Switch	Cisco Nexus N9K-C93180YC-EX	Cisco C93180YC-EX-1	192.168.242.xx
Cisco 10G/25GbE Network Switch	Cisco Nexus N9K-C93180YC-EX	Cisco C93180YC-EX-2	192.168.242.xx
Brocade Fibre Channel SAN Switch	G620	SAN-switch 1	192.168.242.xx
Brocade Fibre Channel SAN Switch	G620	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) virtualized on Oracle Linux Virtualization Manager using Hitachi Virtual Storage Platform E590 and Hitachi Advanced Server DS220.

Database configuration

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

Oracle Database Parameter	Value
Compatible	19.3.0.0.0
Oracle Database size	1 TB
Database storage type	ASM
Database fill factor	80%

Test environment

The following table lists configuration details for VSP E590 testing.

Item	Value
Operating System on VMs	OL 7.9
Workload Type	OLTP/OLAP
Database Size	1 TB
Number of vCPUs	32
Virtual Memory	64 GB per VM
Host Cluster VM Network	2 × 25 Gbps NIC Bonding

Test methodology

The test results were demonstrated using Orion I/O workload and peakmarks OLTP test cases. The peakmarks 9.4 tool was used to validate this solution.

Orion

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

For more information, see *I/O Configuration and Design* in the [Oracle Database Performance Tuning Guide](#).

peakmarks

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

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

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

Test results

The following table provides Orion test result details.

Test Case	Test Result	
	Throughput	RT (ms)
8 KB Random Read	276,073 IOPS	0.96
8 KB Random Write	144,815 IOPS	1.08
1 MB Sequential Read	11.66 GBps	N/A
1 MB Sequential Write	6.62 GBps	N/A

The following table provides peakmarks test result details.

Test Case	Test Results	
	Throughput	RT (ms)
STO-READ	12.159 MBps	N/A
STO-Random 100% read	225,466 IOPS	1.0
SRV-SCAN	46.478 GBps	N/A
DBWR-THR	0.653 GBps	N/A
DA-LOW/ ROW	46.753 GBps	N/A
DA-LOW/STO	27.135 GBps	N/A
TP-LIGHT 1 RPT 0 % update	240,034 TPS	0.4
TP-LIGHT 1 RPT 20 % update	76,134 TPS	1.2
TP-LIGHT 1 RPT 30% update	69,963 TPS	1.3
TP-MEDIUM 25 RPT 0% update	57,006 TPS	1.7
TP-MEDIUM 25 RPT 20% update	2,285 TPS	5.0
TP-MEDIUM 25 RPT 30% update	2,173 TPS	5.4

Conclusion

Hitachi Virtual Storage Platform E590 has been tested and validated as an ideal platform for virtualized workloads. It was tested on Oracle RAC database as well as with multiple environments running at the same time.

Product descriptions

These products are used in this reference architecture.

Hitachi Virtual Storage Platform E series

The Hitachi Virtual Storage Platform E series family provides agile and automated storage built upon the innovative technologies found in our high-end enterprise systems. The expansion of the VSP E series portfolio includes 2 new all-NVMe flash models that deliver super charged, ultra-low latency performance for the business-critical applications that small and midsized businesses rely on.

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 DS220

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

The Intel Xeon Scalable processor family is optimized to address the growing demands on today's IT infrastructure. The server provides 24 slots for high-speed DDR4 memory, allowing up to 3 TB of memory per node when 128 GB DIMMs are used. This server supports up to 12 large form factor storage devices and an additional 2 small form factor storage devices.

This server has three storage configuration options:

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

Hitachi Advanced Server DS120

Optimized for performance, high density, and power efficiency in a dual-processor server, [Hitachi Advanced Server DS120](#) delivers a balance of compute and storage capacity. This 1U rack-mounted server has the flexibility to power a wide range of solutions and applications.

The highly-scalable memory supports up to 3 TB using 24 slots of high-speed DDR4 memory. Advanced Server DS120 is powered by the Intel Xeon Scalable processor family for complex and demanding workloads. There are flexible OCP and PCIe I/O expansion card options available. This server supports up to 12 small form factor storage devices with up to 4 NVMe drives.

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

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

Hitachi Ops Center

Manage, optimize, orchestrate and protect your data with advanced IT analytics and automation using [Hitachi Ops Center](#). Achieve new insights, accelerate resource delivery, eliminate risks, and speed innovation to modernize your data center operations.

Use the power of AI operations with the following:

- [Analyzer](#): Improve IT operations with machine learning (ML) to drive resource service levels, utilization and automation at lower costs. Obtain operational visibility from virtual machines, servers, SAN switches to shared storage resources to optimize an application's full data path.
- [Automator](#): Deliver resources up to 70% faster than manual processes. Free staff to focus on strategic initiatives.
- [Protector](#): Meet tight service level requirements when protecting critical data and applications. Automatically support secondary business functions with data copies staff need to do their jobs. Make better use of backup data for activities, such as e-discovery and analysis. Simplify administration and replication management. Do it all with no disruption to production application availability and performance.

Common Content-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 Storage Area Network switches

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

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

Cisco Nexus switches

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

Hitachi Vantara



Corporate Headquarters

2535 Augustine Drive

Santa Clara, CA 95054 USA

HitachiVantara.com | community.HitachiVantara.com

Contact Information

USA: 1-800-446-0744

Global: 1-858-547-4526

HitachiVantara.com/contact