Hitachi Solution for Databases – Oracle 19c RAC on Oracle Linux KVM Virtualization with Hitachi Advanced Server HA820 G3 and Hitachi Virtual Storage Platform One Block 28

Reference Architecture Guide

MK-SL-302-00
August 2024

Feedback

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

Revision history

Changes	Date
Initial release	August 2024

Reference Architecture Guide

Hitachi Solution for Databases with Oracle Real Application Clusters (RAC) virtualized on an Oracle kernel-based virtual machine (KVM) using a Hitachi Virtual Storage Platform One Block 28 (VSP One Block 28) and Hitachi Advanced Server HA820 G3 with 5th Generation Intel Xeon 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 KVM. In this solution, Oracle RAC Database nodes are deployed on Oracle KVM hosts with Oracle Linux 8.6 UEK 6 and guest virtual machines (VMs) with Oracle Linux 8.10 UEK 6.

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

The Oracle Linux KVM feature provides built-in kernel capabilities to use the Oracle Linux kernel as a hypervisor. It provides a full virtualization solution containing virtualization extensions (Intel VT or AMD-V) on Linux machines to create a two-node Oracle RAC environment.

The following components are used:

- Hitachi Virtual Storage Platform One Block 28 is used as a storage resource.
- Hitachi Advanced Server HA820 G3 with Intel® Xeon® Platinum 8568Y+ Emerald Rapids Processor with 48-core CPUs for storage and computing resources.
- Hitachi Advanced Server HA810 G3 with Intel® Xeon® Gold 6454S Processor 32-core CPUs for the management servers.
- Oracle Linux KVM Hypervisor.
- Oracle Real Application Clusters (RAC).

Running Oracle RAC on Oracle Linux KVM provides the following benefits:

- Oracle Linux KVM is an open source with no licensing cost.
- Conforms to the Oracle hard partitioning licensing requirement.
- Supports full hardware assisted virtualization.

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 One Block 28
- Hitachi Advanced Server HA820 G3 servers
- Hitachi Advanced Server HA810 G3 servers
- Storage Area Networks
- Oracle Database administration
- Oracle Database release 19c, version 19.22
- Oracle Linux 8.6 UEK 6 for KVM hosts and Oracle Linux 8.10 UEK 6 for the Guest OS

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 KVM using Hitachi Virtual Storage Platform One Block 28 and Hitachi Advanced Server HA820 G3. 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 VSP One Block 28.

Business benefits

Oracle Linux KVM consists of a loadable kernel `module kvm.ko`, which provides core virtualization infrastructure and a processor specific module `kvm-intel.ko` or `kvm-armd.ko`.

The following are some benefits of this reference architecture:

- Hard partitioning support complies with Oracle licensing rules.
- Reduces operation and support costs while increasing IT efficiency and agility — on premises and in the cloud with Oracle virtualization.
- Adds or removes servers in the Oracle KVM Manager pool based on business requirements.
- Has virtual machine snapshot capabilities.

- Has online VM migration with VM cloning features.
- Provides a solution for customers who want converged Oracle products — not covered by Hitachi Virtual Storage Platform G700, VSP G800, VSP G900, or VSP 5000 series for Oracle RAC.

High-level infrastructure

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

- Hitachi Advanced Server HA820 G3 servers
- Hitachi Advanced Server HA810 G3 servers
- Hitachi Virtual Storage Platform One Block 28
- Hitachi Virtual Storage Platform E1090
- Brocade G720 64 Gbps SAN infrastructure
- Cisco 10/25 GbE LAN infrastructure

The configuration of Hitachi Virtual Storage Platform One Block 28 and Hitachi Advanced Server HA820 G3 has the following characteristics:

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

Note: The VSP E1090 is optional and not required in your environment. Data volumes can also be configured on the VSP One Block 28 or other VSP storage systems.

The following figure shows the high-level architecture diagram using VSP One Block 28 and HA820 G3 for a Two-Node Oracle 19c RAC configuration.

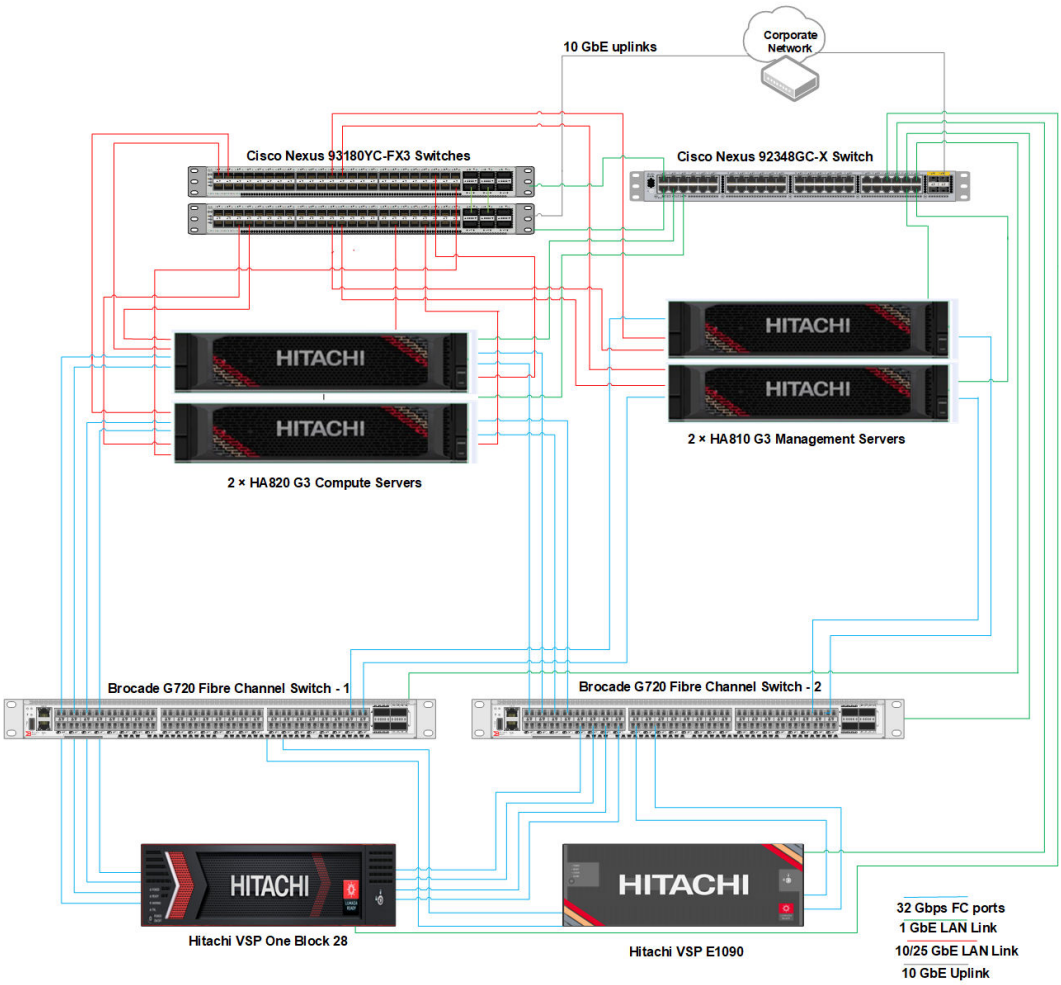

Note: Management servers are not used in this implementation. They are shown for reference only.

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 25 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 21).

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

Vendor	Hardware	Detail Description	Version	Quantity
Hitachi Vantara	Hitachi Virtual Storage Platform One Block 28	3 × CHB pairs 24 × 64 Gbps Fibre Channel ports 1024 GB cache memory 24 × 3.8 TB NVMe SSDs*	A3-02-01-40/00	1
	Hitachi Virtual Storage Platform E1090 – Management node storage system	Two Controllers 8 × 32 Gbps Fibre Channel ports 1024 GB cache memory 48 × 1.9 TB NVMe SSDs*	93-07-01-80/100	1
Hitachi Vantara	Hitachi Advanced Server HA820 G3	2 × Intel Xeon Platinum 8568Y+ 48 Cores CPUs @ 2.3 GHz 64-bit 2 TB memory RDIMM DDR5-5600 MHz	iLO 6: 1.58 Mar 22 2024 System ROM: U54 v2.16 (03/01/2024)	2
		2 × SN1700E 64 Gb 2p FC HBA	Firmware: 14.2.589.19 Driver: lpfc Driver version: 12.8.0.11	
Brocade	G720 Fibre Channel switches	64 × 64 Gbps ports Fibre Channel switch 64 Gbps SFPs	Fabric OS: v9.1.1c	2
Cisco	Cisco Nexus N9K-C93180YC-FX3	48 × 10/25 GbE ports 6 × 40/100 Gbps Quad SFP (QSFP28) ports	BIOS: version 05.51 NXOS: version 10.2(5)	1

Vendor	Hardware	Detail Description	Version	Quantity
	Cisco-C92348GC-X	48 ×1 GbE ports Ethernet switches	BIOS: version 05.51 NXOS: version 10.2(5)	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: Customers can choose larger capacity SSDs to fit their business requirements.

The following table lists software components for compute nodes.

Software	Version	Function
Oracle Linux	8.6 UEK 6with 5.4.17-2136.307.3.1.el8uek.x86_64 on KVM hosts 8.10 UEK 7with 5.4.17-2136.333.5.el8uek.x86_64 on RAC VMs (Guest OS)	Operating system
Oracle Database	19c (Version 19.22.0.0.0)	Database software
Oracle Real Application Clusters	19c (Version 19.22.0.0.0)	Cluster software
Oracle Grid Infrastructure	19c (Version 19.22.0.0.0)	Volume management, file system software, and Oracle automatic storage management
Oracle ASM	2.12	Support ASM storage device for Oracle RAC

Note: According to Oracle DOC ID 1304727.2, (5.4.17-xxxx.el8uek.x86_64) is the minimum supported kernel for OL8.x UEK 6 but we found root.sh failed with CLSRSC-214 and OKA is not supported on this operating system version: '5.4.17-2136.307.3.1.el8uek.x86_64' errors. To resolve this OS and hardware compatibility issue, we upgraded kernel-uek rpm to at least "kernel-uek-5.4.17-2136.316.7.el8uek.x86_64.rpm" or kernel-uek-to 5.4.17-2136.333.5.el8uek" of UEK 6 version.

Solution design

This section describes the reference architecture environment to implement Hitachi Solution for Databases with Oracle RAC virtualized using KVM hypervisor. The environment uses Hitachi Virtual Storage Platform One Block 28.

The infrastructure configuration includes the following:

- Oracle KVM hosts — There are two hosts configured in an Oracle database environment.
- Storage System — There are raw disks 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 systems through Brocade G720 switches.

Storage architecture

This section describes the storage architecture for this solution.

Storage configuration

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

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

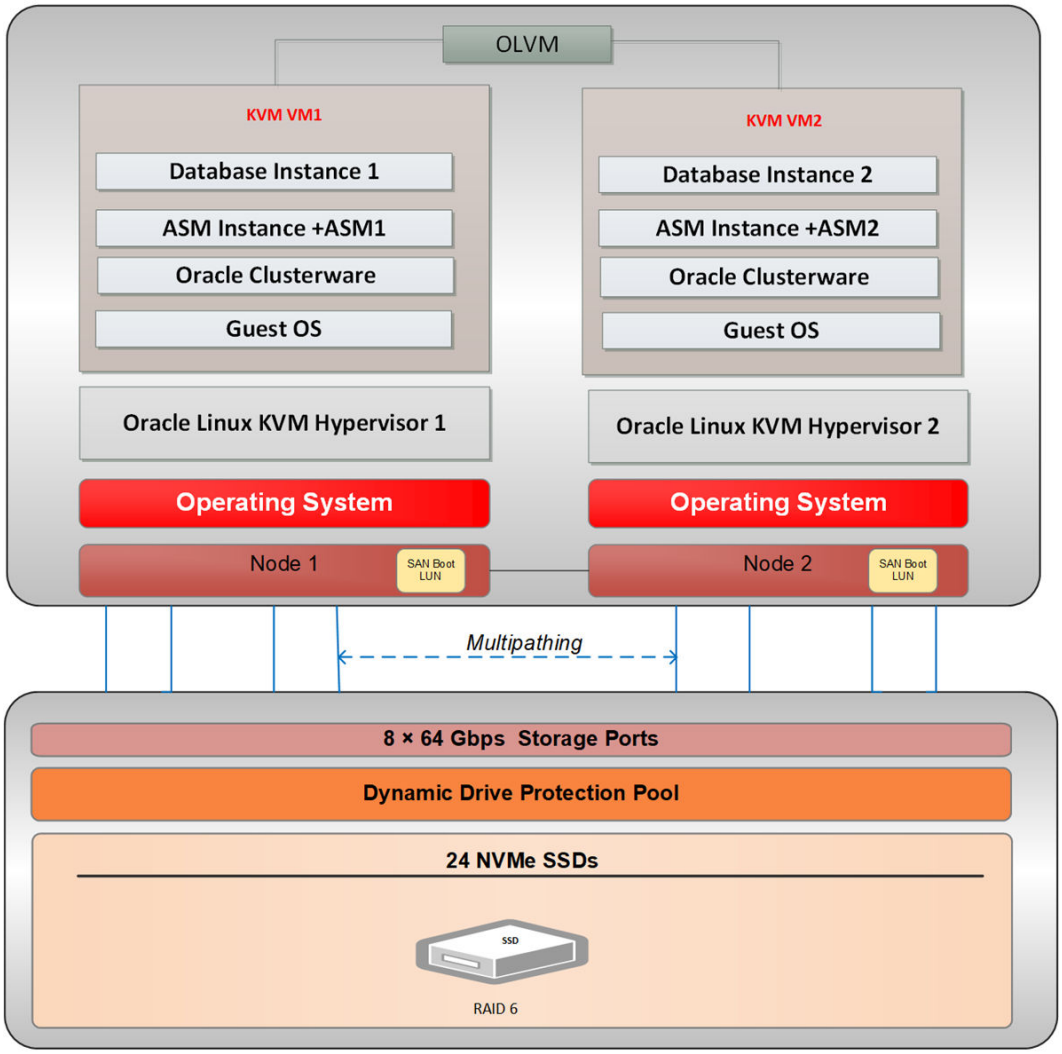

Hitachi VSP One Block 28

The following table shows the VSP One Block 28 storage pool configuration used for this solution.

Pool ID	Oracle RAC-Pool
Pool Type	Dynamic Provisioning Pool
RAID Group	1-1
RAID Level	RAID 6 (6D+2P+1 Spare)

Pool ID	Oracle RAC-Pool
Drive Type	3.8 TB NVMe SSD
Number of Drives	24
LDEV Size	3.8 TB
Pool Capacity	68 TB

In this Solution we used RAID 6 (6D+2P+1Spare) configuration for the storage systems with 24 NVMe SSDs with a single DDP pool. Each SSD Drive capacity is 3.8 TB with DDP enabled, with a total usable capacity of 68 TB.

The following table shows the VSP One Block 28 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	2D, 8D, 1A, 3A, 7A, 5A, 4D, 6D
REDO	6	20 GB	Online Redo Logs and Control Files	
TEMP	3	200 GB	Temp Tablespace	
OCR	3	20 GB	Oracle Cluster Registry and Voting Disk	

The following table shows the VSP E1090 configuration for management servers.

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

Item	Value/Description
Number of Pool LDEVs	4
LDEV Size	3× 1000 GB
Number and Size of CCI Devices	1× 100 MB
Storage Port for Management Servers	7A, 7B, 8A, 8B

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 One Block 28 for small random I/O traffic, such as OLTP transactions. The layout also considers Oracle ASM best practices when using Hitachi Vantara storage.

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

The components in this solution have the flexibility to be used in various deployment scenarios to provide the right balance between performance and ease of management for a given scenario.

Oracle ASM configurations

- Data and Indexes Tablespace — Assign an ASM disk group with external redundancy for the data and index tablespaces.
- TEMP Tablespace — Place the TEMP tablespace in this configuration in the TEMP ASM disk group.
- Online Redo Logs — Create an ASM disk group with external redundancy for Oracle online redo logs.
- Oracle Cluster Registry and Voting Disk — Create an ASM disk group with normal redundancy to contain the OCR and voting disks and to protect against single disk failure to avoid loss of cluster availability. Place each of these files in this configuration in the OCR ASM disk groups.

Oracle initial parameters

The following table lists the Oracle Database settings.

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

Oracle ASM disk mappings

The following table lists 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]1	1 × 350 GB	OS and Oracle Database
OCR	OCR1-OCR3	/dev/xvd[b-d]1	3 × 20 GB	Oracle cluster registry and voting disk
DATA1	DATA1-DATA16	/dev/xvd[e-r]1	14 × 200 GB	Application data
REDO	REDO01-REDO06	/dev/xvd[s-x]1	6 × 20 GB	Online REDO log group
TEMP	TEMP1-TEMP3	/dev/xvd[y-z1]1	3 × 200 GB	Temporary Tablespace

Server and application architecture

This reference architecture uses two Hitachi Advanced Server HA820 G3 servers with 5th Generation Intel Xeon Scalable Processors for each storage system architecture that was tested.

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

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

Server Make and Model	Server Host Name	Role	CPU Type	CPU Core	RAM
Hitachi Advanced Server HA820 G3	KVM server-01	Oracle Linux KVM Hypervisor Host (KVM host1)	2 × Intel Xeon Platinum 8568Y + 48C CPUs	192 (2 × 2 × 48)	2048 GB (64 GB × 2 × 32)
Hitachi Advanced Server HA820 G3	KVM server-02	Oracle Linux KVM Hypervisor Host (KVM host2)	2 × Intel Xeon Platinum 8568Y + 48C CPUs	192 (2 × 2 × 48)	2048 GB (64 GB × 32)
KVM VMs	VM1	Oracle Linux KVM VM for Oracle RAC Node 1	2 × Intel Xeon Platinum 8568Y +	32 vCPU	256 GB
KVM VMs	VM2	Oracle Linux KVM VM for Oracle RAC Node 2	2 × Intel Xeon Platinum 8568Y +	32 vCPU	256 GB

SAN architecture

Map the provisioned LDEVs to multiple ports on Hitachi Virtual Storage Platform One Block 28. These LDEV port assignments provide multiple paths to the storage system from the host for high availability. This reference architecture uses two dual port SN1700E HBAs per Hitachi Advanced Server HA820 G3.

Compute servers

- 8 SAN switch connections are used for Hitachi Virtual Storage Platform One Block 28 Fibre Channel ports.
- 8 SAN switch connections are used for server HBA ports.

Management servers

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

The following table lists details of the Fibre Channel switch connect configuration on Hitachi Virtual Storage Platform One Block 28.

Server	HBA	HBA Port Alias	Switch Zone	Storage System	Storage Port	Brocade G720 Switch
HA820 G3 Server 1	HBA1	CN31_HBA1_1	CN31_HBA1_1_ASE44_204_2D	VSP One Block 28	2D	12(SW1)
	HBA1	CN31_HBA1_2	CN31_HBA1_2_ASE44_204_8D		8D	12(SW2)
	HBA2	CN31_HBA2_1	CN31_HBA2_1_ASE44_204_1A		1A	13(SW1)
	HBA2	CN31_HBA2_2	CN31_HBA2_2_ASE44_204_3A		3A	13(SW2)
HA820 G3 Server 2	HBA1	CN32_HBA1_1	CN32_HBA1_1_ASE44_204_7A	VSP One Block 28	7A	14(SW1)
	HBA1	CN32_HBA1_2	CN32_HBA1_2_ASE44_204_5A		5A	14(SW2)
	HBA2	CN32_HBA2_1	CN32_HBA2_1_ASE44_204_4D		4D	15(SW1)
	HBA2	CN32_HBA2_2	CN32_HBA2_2_ASE44_204_6D		6D	15(SW2)
HA810 G3 Server 1	HBA1	MN33_HBA1_1	MN33_HBA1_1_ASE44_204_1A	VSP E1090	1A	29
	HBA1	MN33_HBA1_2	MN33_HBA1_2_ASE44_204_3A		3A	30
HA810 G3 Server 2	HBA1	MN34_HBA1_1	MN34_HBA1_1_ASE44_204_2A		2A	29
	HBA1	MN34_HBA1_2	MN34_HBA1_2_ASE44_204_4A		4A	30

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 connection.

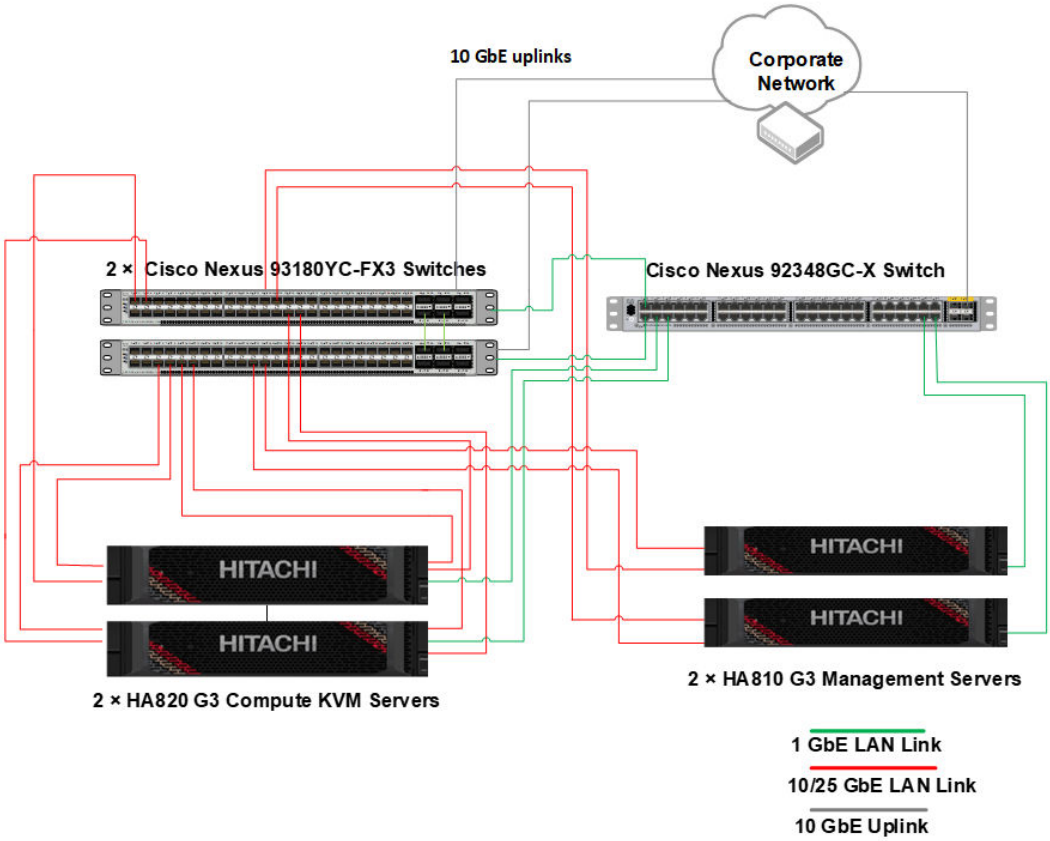

The following table lists the HA820 G3 and HA810 G3 network configuration for this solution.

Server	NIC Ports	VLAN/ Subnet	NIC BOND	IP Address	NW	BW Gbps	Cisco 93180YC-FX3 Switch	
							Switch #	Port
HA820 G3 Server 1	NIC 1-Port 0	100	Bond1	192.168.100.X	Private	10/25	1	32

Server	NIC Ports	VLAN/ Subnet	NIC BOND	IP Address	NW	BW Gbps	Cisco 93180YC-FX3 Switch	
							Switch #	Port
	NIC 2–Port 1					10/25	2	
	NIC 1–Port 1	242	Bond0	192.168.242.X.XXX	Public	10/25	1	33
	NIC 2–Port 0					10/25	2	
	iLO-Dedicated NIC	242	-	192.168.242.X	Private Mgmt	1	-	
HA820 G3 Server 2	NIC 1–Port0	100	Bond1	192.168.100.X	Private	10/25	1	34
	NIC 2–Port 1					10/25	2	
	NIC 1-Port 1	242	Bond0	192.168.242.X	Public	10/25	1	35
	NIC 2–Port 0					10/25	2	
	iLO-Dedicated NIC	242	-	192.168.242.X	Private Mgmt	1	-	
HA810 G3	NIC - 0	242	-	192.168.242.X	Public	10	1	50
	iLO-Dedicated NIC	242	-	192.168.242.X	Private Mgmt	1	-	
HA810 G3	NIC - 0	242	-	192.168.242.X	Public	10	2	50
	iLO-Dedicated NIC	242	-	192.168.242.X	Private Mgmt	1	-	

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

The following table lists the network configuration for servers and VSP One Block 28.

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 One Block 28	192.168.242.xx
VSP One Block 28 CTL1	192.168.242.xx
VSP One Block 28 CTL2	192.168.242xx

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

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

Engineering validation

This section summarizes the key lab verification tests performed on Hitachi Solution for Databases - Oracle Real Application Clusters (RAC) virtualized on Oracle Linux KVM hypervisor using Hitachi Virtual Storage Platform One Block 28 and Hitachi Advanced Server HA820 G3.

Database configuration

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

Oracle Database Parameter	Value
Compatible	19.22.0.0.0
Oracle Database size	2 TB
Database storage type	ASM
Database fill factor	70%

Test environment

The following table lists configuration details for VSP One Block 28 testing.

Item	Value
Operating System on KVM VMs	OL8.6 UEK 6
Workload Type	OLTP/OLAP
Database Size	2 TB
Number of vCPUs	32
Virtual Memory	250 Gb per VM
Host Cluster VM Network	2 × 25 Gbps NIC Bonding

Test methodology

The test results are demonstrated using peakmarks OLTP test cases.

peakmarks is benchmark software for Oracle platforms. It was used in our tests for the following:

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

This provides transparency and comparability in price versus performance considerations for Oracle infrastructures. The peakmarks 10.2 tool was used to validate this solution.

Note: Peakmarks tests conducted on VMs were configured as “Server Class VM” on Oracle Linux KVM Hypervisor.

Conclusion

We performed various database and clusterware validation tests for this unified converged system consisting of a software and hardware stack configured on Hitachi Advanced Server HA820 G3 as a compute node and Hitachi Virtual Storage Platform One Block 28 storage. We found that all the database functionalities for Oracle database are running efficiently without any issues. We validated number of database operations for various OLTP workloads using peakmark and Orion performance benchmarking tools and found that all types of workloads executed smoothly with optimal performance results. .

The Oracle clusterware was configured on High Performance KVM VMs with few CPU and memory resources and achieved the best database performance which is nearly equal to the performance of database clusterware configured on bare metal hosts.

Therefore, we have found that by using the Oracle KVM virtualization feature, customers can get optimal database performance along with saving considerable licensing costs and hardware resources.

Contact your Hitachi sales and engineering team for more details about performance results and best configuration practices.

Product descriptions

These products are used in this reference architecture.

Hitachi Virtual Storage Platform One Block

The Hitachi Virtual Storage Platform One Block series simplifies system setup and management through Hitachi Clear Sight and VSP One Block Administrator. Dynamic Drive Protection reduces RAID complexity, and always-on compression and deduplication enhance simplicity.

Dynamic Carbon Reduction optimizes energy usage by switching CPUs to ECO mode during low activity. Adaptive Data Reduction (ADR) is always on, enhancing efficiency and reducing the overall CO2 footprint.

Thin Image Advanced (TIA) integrates with major snapshot ecosystems, prioritizing security by defending against threats and ensuring data confidentiality. CyberArk Privileged Access Manager plugins enhance block storage system security by prioritizing data confidentiality, ensuring compliance, and actively defending against security threats.

Hitachi Virtual Storage Platform One Block 20 includes 3 dedicated models. All have the same capacity (72 NVMe flash drives, the appliance and 2 × media trays) and they support Fibre Channel, iSCSI, and NVMe TCP connectivity. The new capabilities remove complexity: data reduction is always on, Dynamic Drive Protection removes complicated RAID setup, and Dynamic Carbon Reduction delivers real world reduction in power consumption. In addition, the models are FIPS compliant.

- VSP One Block 24 – 256 GB Cache + SW Advanced Data Reduction (ADR) + 24 cores
- VSP One Block 26 – 768GB Cache + 2x Compression Accelerator Module (CAM) + 24 cores
- VSP One Block 28 – 1TB Cache + 4x CAM + 64 cores

In short, the Hitachi Virtual Storage Platform One Block series combines simplicity, sustainability, and robust security features to optimize system management, energy efficiency, and data protection.

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 HA810 G3

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

Highly scalable memory supports up to 4 TB RAM using 32 slots of 2200 MHz DDR5 RDIMM. HA810 G3 is powered by the Intel Xeon scalable processor family for complex and demanding workloads. There are flexible OCP and PCIe I/O expansion card options available.

The following applications were installed in individual virtual machines in this architecture and would be used in most cases:

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

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

Hitachi Advanced Server HA820 G3

Hitachi Advanced Server HA820 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. Hitachi Advanced Server HA820 G3 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, HA800 G3 series 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 supports up to 8 TB RAM using 32 slots of 4800 MHz DDR5 RDIMM. HA820 G3 is powered by the Intel Xeon scalable 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.

Oracle Linux KVM (OL-KVM)

Oracle Linux KVM is a feature of Oracle Linux. With the Unbreakable Enterprise Kernel (UEK) Release 6, the Oracle Linux server virtualization solution with KVM has been enhanced. Users can take either a previously deployed version of Oracle Linux and turn the OS into a KVM host, or a KVM configuration can be set up from a base Oracle Linux installation. Oracle Linux KVM is the same hypervisor used in Oracle Cloud Infrastructure, giving users an easy migration path to move workloads into Oracle Cloud.

There are two types of VMs configured for Hitachi solutions for Oracle database namely VMs optimized for *Server class* and VMs optimized for *High Performance VMs*.

With *High Performance VMs* we can configure a virtual machine for high performance, so that it runs with performance metrics as close to bare metal as possible. When you choose high performance optimization, the virtual machine is configured with a set of automatic and recommended manual settings for maximum efficiency.

- The high-performance option is supported by Red Hat Virtualization 4.2 and later
- To change the optimization mode of a new or existing virtual machine to high performance, we may need to make manual changes to the cluster and to the pinned host configuration first.

Oracle Linux

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

Oracle Database with Real Application Clusters Option

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

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

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

VMware ESXi

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

VMware vCenter Server Appliance

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

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

Brocade switches from Broadcom

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

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

Cisco Nexus switches

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

Peakmarks

Test 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.

Hitachi Vantara

Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA

HitachiVantara.com/contact