

WHITE PAPER

# Hitachi Solution for Databases - Oracle Real Application Clusters Database 12c based on Hitachi Advanced Server DS220 and Virtual Storage Platform G900

Reference Architecture Guide

By Amol Bhoite

June 2018

# Feedback

Hitachi Vantara welcomes your feedback. Please share your thoughts by sending an email message to [SolutionLab@HitachiVantara.com](mailto: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

Revision	Changes	Date
MK-SL-062-00	Initial release.	May 16, 2018
MK-SL-062-01	Update the title of the paper to reflect key solution components.	May 31, 2018

# Table of Contents

<b>Solution Overview</b>	<b>2</b>
Business Benefits	2
High Level Infrastructure	2
<b>Key Solution Components</b>	<b>4</b>
Hitachi Virtual Storage Platform F Series Family	7
Hitachi Advanced Server DS220 Server	7
Hitachi Advanced Server DS120 Server	8
Red Hat Enterprise Linux	8
Device Mapper Multipathing	8
Oracle Database With the Real Application Clusters Option	8
Hitachi Infrastructure Analytics Advisor	8
Hitachi Storage Advisor	9
Oracle Enterprise Manager	9
Hitachi Storage Adapter for Oracle Enterprise Manager	9
Hitachi Server Adapter for Oracle Enterprise Manager	9
Brocade Switches	10
Cisco Switches	10
<b>Solution Design</b>	<b>10</b>
Storage Architecture	10
Server and Application Architecture	16
SAN Architecture	17
Network Architecture	19
Hitachi Applications	22
<b>Engineering Validation</b>	<b>25</b>
Test Methodology	25
Test Results	26

# Hitachi Solution for Databases - Oracle Real Application Clusters Database 12c based on Hitachi Advanced Server DS220 and Virtual Storage Platform G900

## Reference Architecture Guide

Use this reference architecture guide to see how Hitachi Solution for Databases provides a high performance, integrated, converged solution for Oracle. The environment uses Hitachi Virtual Storage Platform G900 (VSP G900), Hitachi Advanced Server DS220 with Intel Xeon Gold 6140 processors, and Hitachi Advanced Server DS120 Servers with Intel Xeon Silver 4110 processors. With these products, design an Oracle converged infrastructure to meet your requirements and budget.

This Hitachi Unified Compute Platform CI architecture for Oracle Database is engineered, pre-tested, and qualified to provide predictable performance and the highest reliability in demanding, dynamic Oracle environments. This solution is validated to ensure consistent, predictable results.

This proven solution to optimize your Oracle database environment integrates servers, storage systems, and networks. The environment provides reliability, high availability, scalability, and performance while processing small-scale to large-scale online transaction processing (OLTP) and online analytical processing (OLAP) workloads. The dedicated servers run Oracle Database 12c R1 with the Oracle Real Application Clusters (RAC) option. The operating system is Red Hat Enterprise Linux 7.4. This reference architecture document is for you if you are in one of the following roles:

- Database administrator
- Storage administrator
- Database performance analyzer
- IT professional with the responsibility of planning and deploying an Oracle Database solution

To use this reference architecture guide, you need familiarity with the following:

- Hitachi Virtual Storage Platform G900
- Hitachi Advanced Server DS220 servers
- Hitachi Advanced Server DS120 servers
- Storage area networks
- Oracle RAC Database 12c Release 1
- Oracle Automatic Storage Management (Oracle ASM)
- Hitachi Adapters for Oracle Database
- Hitachi Storage Adapter for Oracle Enterprise Manager
- Hitachi Server Adapter for Oracle Enterprise Manager
- Red Hat Enterprise Linux
- Red Hat Enterprise Linux Device-Mapper Multipath

---

**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

This reference architecture implements Hitachi Unified Compute Platform CI for Oracle Real Application Clusters on four nodes using Hitachi Virtual Storage Platform G900. This environment addresses the high availability, performance, and scalability requirements for OLTP and OLAP workloads. Your solution implementation can be tailored to meet your specific needs.

To keep up with the growth of Oracle databases, companies have rushed to add storage and servers without considering how existing storage resources might be better utilized. This often requires the addition of more administrators and software layers to manage the expansion of resources, and the purchase of more Oracle DB licenses with a corresponding increase in support costs.

Without a holistic strategy for optimizing the Oracle environment, unsystematic growth and its associated uncontrolled costs can create an increasingly expensive cycle. Through a combination of our solutions and expertise in Oracle environments, Hitachi Vantara can help companies control capital and operational costs, while increasing performance, data protection, and flexibility. We can help ensure the right information is available in the right place, at the right time – and for the right cost.

## Business Benefits

This reference architecture provides the following benefits:

- Predictable, repeatable, reliable results that are pre-validated
- Faster speed to deploy and increased ability to meet changing needs, with a single source for components and prescriptive guides
- Extreme reliability of Hitachi Storage
- Simplified and centralized storage management
- More effective and efficient use of Oracle Database licensing

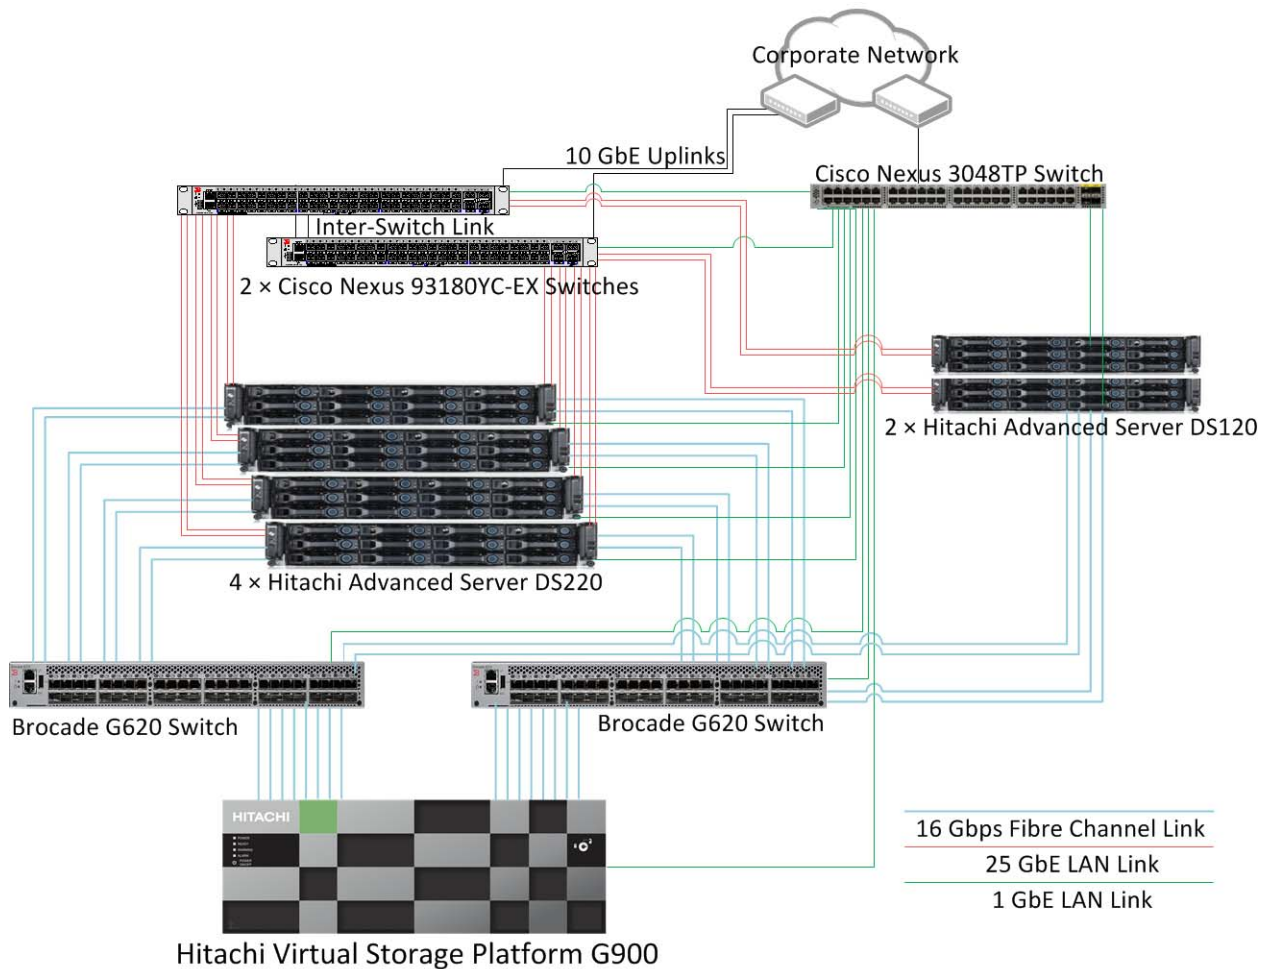
## High Level Infrastructure

Figure 1 shows the high-level infrastructure for this solution.

The configuration of Virtual Storage Platform G900 and Hitachi Advanced Server DS220 have the following characteristics:

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

Figure 1



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

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

Uplink speed to the corporate network will be dependent on the customer environment and requirements. The Cisco Nexus 93180YC-EX switches can support uplink speeds of 25 GbE, 40 GbE or 100 GbE if higher bandwidth is required.

## Key Solution Components

The key solution components for this solution are listed in Table 1, "Hardware Components," on page 4 and Table 2, "Software Components," on page 6.

TABLE 1. HARDWARE COMPONENTS

Hardware	Detailed Description	Firmware/Version	Quantity
Hitachi Virtual Storage Platform G900	<ul style="list-style-type: none"> <li>■ Two Controllers</li> <li>■ 20 x 16 Gbps Fibre Channel Ports</li> <li>■ 8 x 12 Gbps Backend SAS Ports</li> <li>■ 512 GB cache memory</li> <li>■ 64 x 1.9 TB SSDs Plus 2 spares</li> <li>■ 4 x 3 TB 7.2krpm SAS drives</li> </ul>	88-01-03-60/00	1
Hitachi Advanced Server DS220 servers	<ul style="list-style-type: none"> <li>■ 2 x Intel Xeon Gold 6140 CPU @ 2.30GHz</li> <li>■ 768 GB (64 GB x 12) DIMM DDR4 Synchronous Registered (Buffered) 2666 MHz</li> </ul>	BIOS:3A10.H3 BMC: 3.75.06 CPLD:10	4
	<ul style="list-style-type: none"> <li>■ 2 x Intel Corporation Ethernet Controller XXV710 for 25 GbE SFP28</li> </ul>	<ul style="list-style-type: none"> <li>■ Driver: i40e</li> <li>■ Version: 2.1.26</li> <li>■ Firmware: 5.51 0x80002bca 1.1568.0</li> </ul>	
	<ul style="list-style-type: none"> <li>■ 2 x Emulex LightPulse LPe31002-M6 2-Port 16 Gb Fibre Channel Adapter</li> </ul>	Boot: 11215621 Firmware: 11.2.156.27	

TABLE 1. HARDWARE COMPONENTS (CONTINUED)

Hardware	Detailed Description	Firmware/Version	Quantity
Hitachi Advanced Server DS120 servers	<ul style="list-style-type: none"> <li>■ 2 x Intel Xeon Silver Processor 4110, 8-core, 2.1GHz, 85W</li> <li>■ 8 x 32GB DDR4 R-DIMM 2666Mhz (256GB total)</li> <li>■ 1 x 64GB SATADOM</li> </ul>	BIOS:3A10.H3 BMC: 3.75.06 CPLD:10	2
	<ul style="list-style-type: none"> <li>■ 1 x Intel Corporation Ethernet Controller XXV710 for 25 GbE SFP28</li> </ul>	<ul style="list-style-type: none"> <li>■ Driver: i40e</li> <li>■ Version: 2.1.26</li> </ul> Firmware: 5.51 0x80002bca 1.1568.0	
	<ul style="list-style-type: none"> <li>■ 1 x Emulex LightPulse LPe31002-M6 2-Port 16Gb Fibre Channel Adapter</li> </ul>	Boot: 11215621 Firmware: 11.2.156.27	
Brocade Fibre Channel Switches	<ul style="list-style-type: none"> <li>■ G620</li> <li>■ 48 port Fibre Channel switch</li> <li>■ 16 Gbps SFPs</li> <li>■ Brocade hot-pluggable SFP+, LC connector</li> </ul>	Kernel: 2.6.34.6 Fabric OS: v8.0.1	2
Cisco Nexus	<ul style="list-style-type: none"> <li>■ 93180YC-EX</li> <li>■ 48 x 10/25 GbE fiber ports</li> <li>■ 6 x 40/100 Gbps Quad SFP (QSFP28) ports</li> </ul>	BIOS: version 07.61 NXOS: version 7.0(3)I4(7)	2
Cisco Nexus	<ul style="list-style-type: none"> <li>■ 3048TP</li> <li>■ 1 GE 48-Port Gb Ethernet Switch</li> </ul>	BIOS: version 4.0.0 NXOS: version 7.0(3)I4(7)	2



TABLE 2. SOFTWARE COMPONENTS

Software	Version	Function
Red Hat Enterprise Linux	RHEL 7.4 (Kernel Version - 3.10.0-693.21.1.el7.x86_64)	Operating System
Oracle 12c	12c Release 1 (12.1.0.2.0)	Database Software
Oracle Real Application Cluster	12c Release 1 (12.1.0.2.0)	Cluster Software
Oracle Grid Infrastructure	12c Release 1 (12.1.0.2.0)	Volume Management, File System Software, and Oracle Automatic Storage Management
Red Hat Enterprise Linux Device Mapper Multipath	-	Multipath Software
Hitachi Storage Navigator	Microcode dependent	Storage management Software
VMware ESXi	Version 6.5.0.13000 Build 7515524	ESXi for management nodes
VMware vCenter Server	Version 6.5.0 build 4602587	Management cluster
Hitachi Storage Advisor (HSA)	2.3	Storage orchestration software
Hitachi Infrastructure Analytics Advisor (HIAA)	3.3	Analytics Software
Manager for Hitachi adapters for Oracle Database	2.2.0	Hitachi adapters management Virtual appliance software
Hitachi Storage Adapter for Oracle Enterprise Manager	2.2.0	Storage management software
Hitachi Server Adapter for Oracle Enterprise Manager	2.2.0	Server management software
Oracle Enterprise Manager Cloud Control 13c	13c Release 2 (13.2.0.0)	OEM software
Oracle Enterprise Manager Cloud Control 13c plug-ins	13c Release 2	Hitachi Storage and Server OEM plugins

TABLE 3. SOFTWARE COMPONENTS FOR MANAGEMENT NODES

Software	Version	Function
VMware ESXi	Version 6.5.0.13000 Build 7515524	ESXi for management nodes
VMware vCenter Server	Version 6.5.0 build 4602587	Management cluster
Hitachi Storage Advisor (HSA)	2.3	Storage orchestration software
Hitachi Infrastructure Analytics Advisor (HIAA)	3.3	Analytics Software
Manager for Hitachi adapters for Oracle Database	2.2.0	Hitachi adapters management Virtual appliance software
Hitachi Storage Adapter for Oracle Enterprise Manager	2.2.0	Storage management software
Hitachi Server Adapter for Oracle Enterprise Manager	2.2.0	Server management software
Oracle Enterprise Manager Cloud Control 13c	13c Release 2 (13.2.0.0)	OEM software
Oracle Enterprise Manager Cloud Control 13c plug-ins	13c Release 2	Hitachi Storage and Server OEM plugins
Virtual SVP (vSVP)	Microcode dependent	Storage management software

### Hitachi Virtual Storage Platform F Series Family

Use [Hitachi Virtual Storage Platform F series family](#) storage for a flash-powered cloud platform for your mission critical applications. This storage meets demanding performance and uptime business needs. Extremely scalable, its 4.8 million random read IOPS allows you to consolidate more applications for more cost savings.

Hitachi Storage Virtualization Operating System RF is at the heart of the Virtual Storage Platform F series family. It provides storage virtualization, high availability, flash optimized performance, quality of service controls, and advanced data protection. This proven, mature software provides common features, management, and interoperability across the Hitachi portfolio. This means you can reduce migration efforts, consolidate assets, reclaim space, and extend life.

Reduce risks and solve problems faster. Integrated power analytics and automation features bring artificial intelligence to your data center. Cloud-accessible monitoring tools give your product support experts access wherever they have an internet connection for fast troubleshooting and remediation.

This solution uses Virtual Storage Platform F900/G900, which supports [Oracle Real Application Clusters](#).

### Hitachi Advanced Server DS220 Server

[Hitachi Advanced Server DS220](#) is a general-purpose 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 DS220 provides flexible I/O scalability for today's diverse data center application requirements.

## Hitachi Advanced Server DS120 Server

[Hitachi Advanced Server DS120](#) provides flexible and scalable configurations for hyper-converged datacenters, provides computing performance, sophisticated power and thermal design to avoid unnecessary OPEX with quick deployment. For this solution two DS120 servers are used. The two DS120 servers are configured as a VMware vCenter cluster. Virtual machines on the cluster are used to host management applications. The management applications installed depend on customer needs and requirements. The following applications were installed in individual virtual machines in this architecture and would be installed in most cases.

- vCenter
- Oracle Enterprise Manager (OEM) 13c
- Oracle Adapter Manager
- Hitachi Storage Advisor (HSA)
- Hitachi Infrastructure Analytics Advisor / Hitachi Datacenter Analytics (HIAA/HDCA)
- HDCA Probe

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

## Red Hat Enterprise Linux

[Red Hat Enterprise Linux](#) delivers military-grade security, 99.999% uptime, support for business-critical workloads, and so much more. Ultimately, the platform helps you reallocate resources from maintaining the status quo to tackling new challenges.

## Device Mapper Multipathing

[Device mapper multipathing](#) (DM-Multipath) allows you to configure multiple I/O paths between server nodes and storage arrays into a single device.

These I/O paths are physical SAN connections that can include separate cables, switches, and controllers. Multipathing aggregates the I/O paths, creating a new device that consists of the aggregated paths.

## Oracle Database With the Real Application Clusters Option

[Oracle Database](#) has a multi-tenant architecture so you can 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 a file system for Oracle database files. This supports 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.

## Hitachi Infrastructure Analytics Advisor

With [Hitachi Infrastructure Analytics Advisor](#), you can define and monitor storage service level objectives (SLOs) for resource performance. You can identify and analyze historical performance trends to optimize storage system performance and plan for capacity growth.

Use Hitachi Infrastructure Analytics Advisor to register resources (storage systems, hosts, servers, and volumes), and set service-level thresholds. You are alerted to threshold violations and possible performance problems (bottlenecks). Using analytics tools, you find which resource has a problem and analyze its cause to help solve the problem. The Infrastructure Analytics Advisor ensures the performance of your storage environment based on real-time SLOs.

## Hitachi Storage Advisor

[Hitachi Storage Advisor](#) is an infrastructure management solution that unifies storage management solutions such as storage provisioning, data protection, and storage management; simplifies the management of large scale data centers by providing smarter software services; and is extensible to provide better programmability and better control.

## Oracle Enterprise Manager

[Oracle Enterprise Manager](#) provides a “single pane of glass” that allows you to manage on-premises and cloud-based IT using the same familiar interface you know and use on-premises every day. Oracle Enterprise Manager today is the nerve center of IT operations among thousands of enterprises. Millions of assets in Oracle’s SaaS and PaaS public cloud operations are managed by Enterprise Manager round the clock.

Enterprise Manager is the industry’s first complete cloud solution with [Cloud Management](#). This includes self-service provisioning balanced against centralized, policy-based resource management, integrated chargeback and capacity planning, and complete visibility of the physical and virtual environments from applications to disk.

This solution uses Oracle Enterprise Manager Cloud Control, version 13c release 2. This allows you to use these cloud management features:

- Use the Database Cloud Self Service Portal
- Benefit from the Improved Service Catalog
- Perform Snap Cloning using “Test Master Snapshot”
- Take advantage of the Chargeback and Consolidation Planner plugins

For more information, see New Features in [Oracle Enterprise Manager Cloud Control 13c](#)

## Hitachi Storage Adapter for Oracle Enterprise Manager

[Hitachi Storage Adapter](#) for Oracle Enterprise Manager presents an integrated, detailed view of the Hitachi storage supporting your Oracle databases. By gaining visibility into capacity, performance and configuration information, administrators can manage service levels more effectively, and ensure service level agreements (SLAs) are met to support business goals.

## Hitachi Server Adapter for Oracle Enterprise Manager

[Hitachi Server Adapter](#) for Oracle Enterprise Manager is an Oracle Enterprise Manager plug-in that enables monitoring of Hitachi Advanced servers in Oracle Enterprise Manager.

For Hitachi Advanced servers, it provides visibility into the components, including their status, health, and attributes. In addition, the adapter supplies information about any Oracle database instances running on the servers. Both RAC and non-RAC databases are supported.

## *VMware Esxi*

[VMware ESXi](#) is the next-generation hypervisor, providing a new foundation for virtual infrastructure. This innovative architecture operates independently from any general-purpose operating system, offering improved security, increased reliability, and simplified management.

## *vCenter Appliance*

[The 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 comes as 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, including vCenter SSO (Single Sign-on), Inventory Service, vSphere Web Client and the vCenter Server itself.

## **Brocade Switches**

Brocade and Hitachi Vantara partner 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.

SAN switches are optional and direct connect is also possible under certain circumstances, but customers should check the support matrix to ensure support prior to implementation.

The solution uses the following Brocade products:

- Brocade G620, 48 port Fibre Channel

## **Cisco Switches**

The Cisco Nexus Switch product line provides a series of solutions that can 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.

The solution uses the following Cisco products:

- Nexus 93180YC-EX, 48-port 10/25 GbE switch
- Nexus 3048TP, 48-port 1GbE Switch

## **Solution Design**

This describes the reference architecture environment to implement Hitachi Unified Compute Platform CI for Oracle with the Real Application Clusters option. The environment uses Hitachi Virtual Storage Platform G900.

The infrastructure configuration includes the following:

- **Oracle RAC Servers** - There are four server nodes configured in an Oracle Real Application Cluster.
- **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 G620 switches. In this solution SAN switches are optional. Direct connect is possible under certain circumstances. Check the support matrix to ensure support for your choice.

## **Storage Architecture**

This describes the storage architecture for this solution.

### *Storage Configuration*

It takes into consideration Hitachi Vantara for Hitachi Virtual Storage Platform and Oracle recommended practices for the design and deployment of database storage.

The high-level storage configuration diagram for this solution is shown in Figure 2.

Figure 2

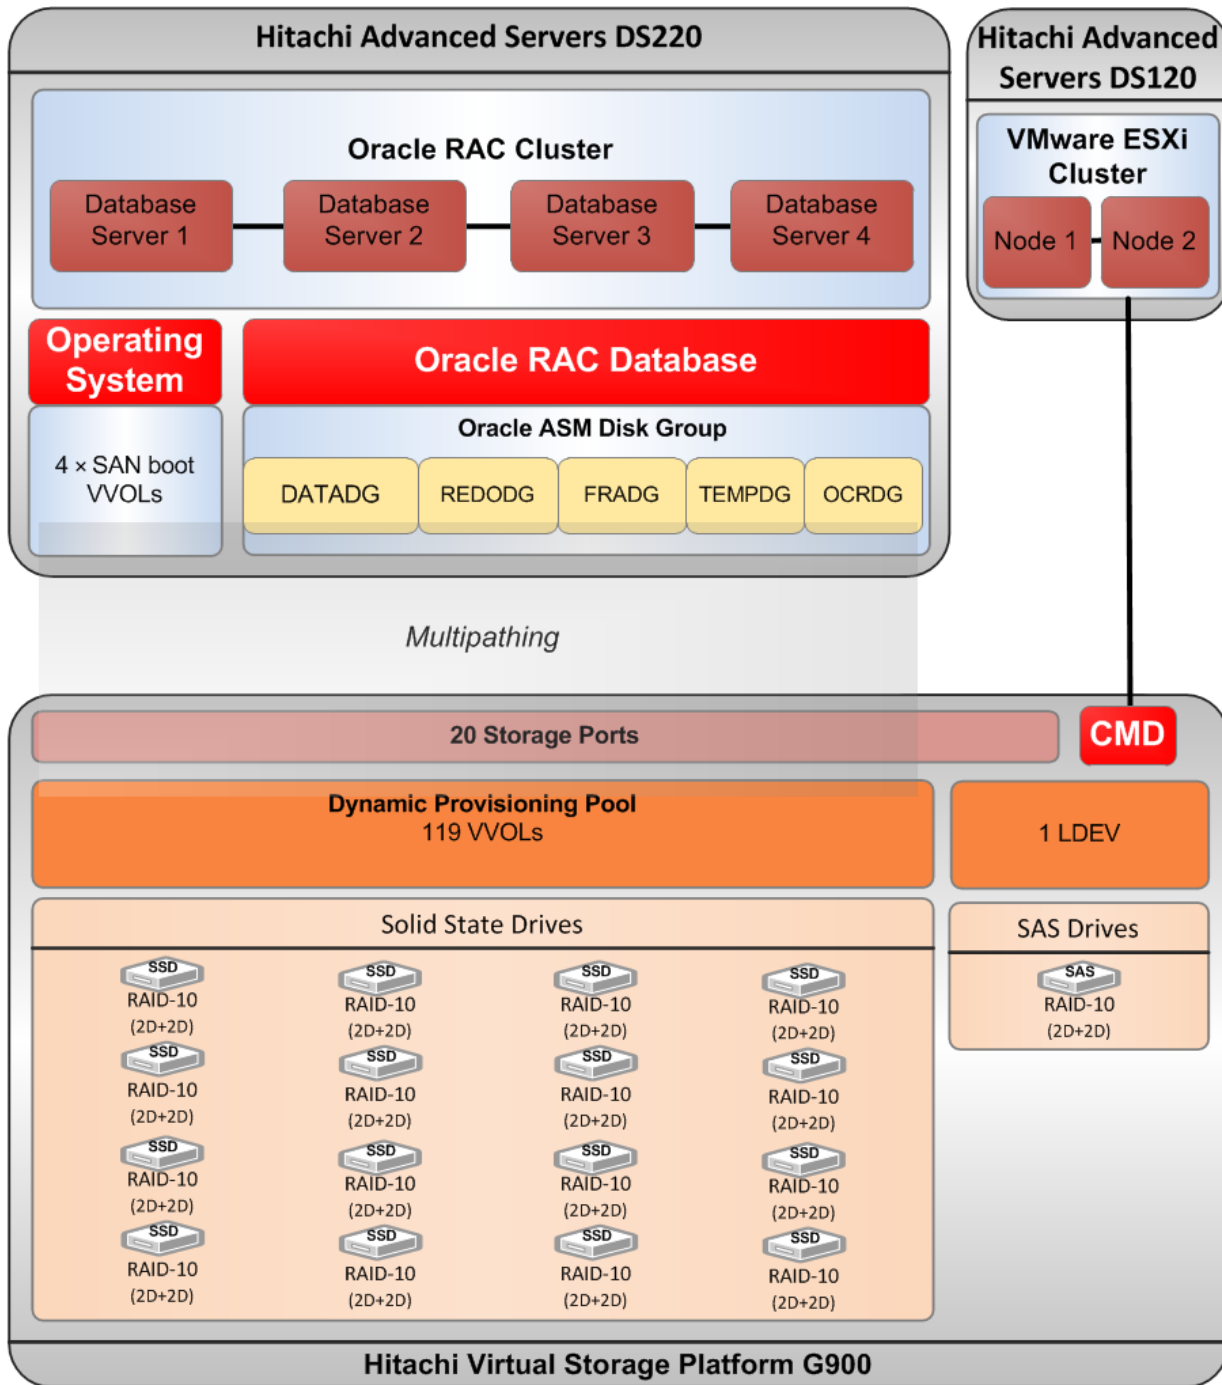


Table 4 shows the storage pool configuration used for this solution.

TABLE 4. STORAGE POOL CONFIGURATION

Pool ID	Hitachi-UCP-CI-Oracle
Pool Type	Dynamic Provisioning
RAID Group	1-1 – 1-16
RAID Level	RAID-10 (2D+2D)
Drive Type	1.9 TB SSDs
Number of Drives	64
Number of Spare Drives	2
Number of Pool Volume LDEVs	64
Pool Volume LDEV size	880 GB
Pool Capacity	54.99 TB

Table 5 shows the logical storage configuration used in this solution.

TABLE 5. LOGICAL STORAGE CONFIGURATION

Pool ID	Hitachi-UCP-CI-Oracle
Number of VVOLs	119
VVOL Size	4 × 200 GB, 64 × 160 GB, 3 × 5 GB, 16 × 40 GB, 16 × 10 GB, 16 × 40 GB
Purpose	<ul style="list-style-type: none"> <li>■ Operating System</li> <li>■ Oracle <ul style="list-style-type: none"> <li>■ System</li> <li>■ Sysaux</li> <li>■ Undo</li> <li>■ Temp</li> <li>■ Redo Logs</li> <li>■ Parameter and Password file</li> <li>■ Oracle Cluster Registry and Voting Disk</li> </ul> </li> </ul>
Storage Port	1A, 2A, 3A, 4A, 5A, 6A, 7A, 8A, 1B, 2B, 3B, 4B, 5B, 6B, 7B, 8B, 1C,1D,2C,2D

There is an additional RAID group consisting of four 3 TB 7.2krpm SAS drives configured as RAID-10 (2D+2D). This is used as shared storage for the management server cluster. A single 3 TB LUN is mapped to four storage ports. Additional LUNs can be mapped if required. While the test environment was configured using a dedicated SAS RAID group for the management server cluster, this can be configured as a dedicated SSD RAID group, a dedicated HDP pool, or it can use capacity on the HDP pool configured for the Oracle environment depending on customer requirements.

### Database Layout

The database layout design uses recommended practices from Hitachi Vantara for Hitachi Virtual Storage Platform G900 for small random I/O traffic, such as OLTP transactions. The layout also takes into account the Oracle ASM best practices when using Hitachi storage. Base the storage design for database layout needs on the requirements of a specific application implementation. The design can vary greatly from one implementation to another based on the RAID configuration and number of drives used during the implementation. The components in this solution set 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 Configuration

- **Data and Indexes Tablespace** — Assign an ASM diskgroup with external redundancy for the data and index tablespaces.
- **TEMP Tablespace** — Place TEMP tablespace in this configuration in the Data ASM diskgroup.
- **Undo Tablespace** — Create an UNDO tablespace in this configuration within the Oracle Data ASM diskgroup. Assign one UNDO tablespace for each node in the Oracle RAC environment.
- **Online Redo Logs** — Create 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.
- **Database Block Size Settings** — Set the database block size to 8 KB.
- **ASM FILE SYSTEM I/O Settings** — Set the Oracle ASM I/O operations for database files, as follows:
  - FILESYSTEMIO\_OPTIONS = setall

Table 6 shows the Oracle RAC Database Settings.

TABLE 6. ORACLE RAC DATABASE SETTINGS

Environment	Value
RAC configuration	Yes
ASM	Yes - Oracle RAC Database

Table 7 shows the Oracle Environment Parameters.

TABLE 7. ORACLE ENVIRONMENT PARAMETERS

Setting	Value
DB_CLOCK_SIZE	8 KB
SGA_TARGET	400 GB
PGA_AGGREGATE_TARGET	192 GB
DB_CACHE_SIZE	172 GB
DB_KEEP_CACHE_SIZE	96 GB
DB_RECYCLE_CACHE_SIZE	24 GB
INMEMORY_SIZE	48 GB
USE_LARGE_PAGES	TRUE
FILESYSTEMIO_OPTIONS	SETALL
DISK_ASYNCH_IO	TRUE

Table 8 shows the details of the disk mappings from the LUNs to the ASM disk groups for Oracle RAC Database tablespaces.

TABLE 8. LUNS AND ORACLE ASM DISK MAPPINGS

ASM Disk Group	ASM Disk	DM-Multipath LUNs	LUN Details	Purpose
NA	NA	/dev/mapper/mpatha	4 × 200 GB	OS and Oracle four node RAC Database
OCRDG	OCRDISK1 - OCRDISK3	/dev/mapper/mpathaa - /dev/mapper/mpathac	3 × 5 GB	Oracle Cluster Registry and Voting Disk
REDODG	REDODISK1 - REDODISK16	/dev/mapper/mpathad - /dev/mapper/mpathap /dev/mapper/mpathba - /dev/mapper/mpathbc	16 × 10 GB	Online REDO log group
FRADG	FRADISK1 - FRADISK16	/dev/mapper/mpathbd - /dev/mapper/mpathbp /dev/mapper/mpathcb - /dev/mapper/mpathcc	16 × 40 GB	Flash Recovery Area

TABLE 8. LUNS AND ORACLE ASM DISK MAPPINGS (CONTINUED)

ASM Disk Group	ASM Disk	DM-Multipath LUNs	LUN Details	Purpose
TEMPDG	TEMPDISK1 - TEMPDISK16	/dev/mapper/mpathcd - /dev/mapper/mpathcp /dev/mapper/mpathda - /dev/mapper/mpathdc	16 × 40 GB	Temp
DATADG	DATADISK1 - DATADISK64	/dev/mapper/mpathdd - /dev/mapper/mpathdp /dev/mapper/mpatheb - /dev/mapper/mpathep /dev/mapper/mpathfa - /dev/mapper/mpathfp /dev/mapper/mpathga - /dev/mapper/mpathgp	64 × 160 GB	Application Data

### Server and Application Architecture

This reference architecture uses four Hitachi Advanced Server DS220 servers for a four-node Oracle RAC configuration.

This provides the compute power for the Oracle RAC database to handle complex database queries and a large volume of transaction processing in parallel. Table 9 describes the details of the server configuration for this solution.

This reference architecture uses two Hitachi Advanced Server DS120 servers for *VMware ESXi* management server configuration.

Details of the VMware ESXi management servers are specified in Table 9.

TABLE 9. HITACHI ADVANCED SERVER DS220 AND DS120 SERVER SPECIFICATIONS

Hitachi Advanced Server	Server	Server Name	Role	CPU Core	RAM
DS220	<b>Oracle Server1</b>	oracle-rac-01	Oracle RAC node 1	36	768 GB (64 GB × 12)
	<b>Oracle Server2</b>	oracle-rac-02	Oracle RAC node 2	36	768 GB (64 GB × 12)
	<b>Oracle Server3</b>	oracle-rac-03	Oracle RAC node 3	36	768 GB (64 GB × 12)
	<b>Oracle Server4</b>	oracle-rac-04	Oracle RAC node 4	36	768 GB (64 GB × 12)

TABLE 9. HITACHI ADVANCED SERVER DS220 AND DS120 SERVER SPECIFICATIONS (CONTINUED)

Hitachi Advanced Server	Server	Server Name	Role	CPU Core	RAM
DS120	Management server	VMware ESXi 1	Hitachi Storage Advisor VM	16	256 GB (32 GB × 8)
		VMware ESXi 2	Hitachi Infrastructure Analytics Advisor VM Manager for Hitachi Adapters for Oracle Database VM Oracle Enterprise Manager Cloud Control 13c VM	16	256 GB (32 GB × 8)

### SAN Architecture

Map the provisioned LDEVs to multiple ports on Hitachi Virtual Storage Platform G900 (VSP G900). These LDEV port assignments provide multiple paths to the storage system from the host for high availability.

- 20 SAN switch connections are being used for VSP G900 host ports.
- 20 SAN switch connections are being used for server HBA ports.

Table 10 shows details of the Fibre Channel switch connect configuration on the Hitachi Virtual Storage Platform G900 ports.

TABLE 10. SAN HBA CONNECTION CONFIGURATION TO VSP G900

Server	HBA Ports	Storage Host Group	Switch Zone	Storage System	Storage Port	Brocade G620 Switch
DS220 Server 1	HBA1_1	CN1_HBA1_1	CN1_HBA1_1_G900_1A	VSP G900	1A	69
	HBA1_2	CN1_HBA1_2	CN1_HBA1_2_G900_1B		1B	70
	HBA2_1	CN1_HBA2_1	CN1_HBA2_1_G900_2A		2A	69
	HBA2_2	CN1_HBA2_2	CN1_HBA2_2_G900_2B		2B	70
DS220 Server 2	HBA1_1	CN2_HBA1_1	CN2_HBA1_1_G900_3A		3A	69
	HBA1_2	CN2_HBA1_2	CN2_HBA1_2_G900_3B		3B	70
	HBA2_1	CN2_HBA2_1	CN2_HBA2_1_G900_4A		4A	69
	HBA2_2	CN2_HBA2_2	CN2_HBA2_2_G900_4B		4B	70
DS220 Server 3	HBA1_1	CN3_HBA1_1	CN3_HBA1_1_G900_5A		5A	69
	HBA1_2	CN3_HBA1_2	CN3_HBA1_2_G900_5B		5B	70
	HBA2_1	CN3_HBA2_1	CN3_HBA2_1_G900_6A		6A	69
	HBA2_2	CN3_HBA2_2	CN3_HBA2_2_G900_6B		6B	70
DS220 Server 4	HBA1_1	CN4_HBA1_1	CN4_HBA1_1_G900_7A		7A	69
	HBA1_2	CN4_HBA1_2	CN4_HBA1_2_G900_7B		7B	70
	HBA2_1	CN4_HBA2_1	CN4_HBA2_1_G900_8A		8A	69
	HBA2_2	CN4_HBA2_2	CN4_HBA2_2_G900_8B		8B	70
DS120 Server 1	HBA1_1	MN1_HBA1_1	MN1_HBA1_1_G900_1C	1C	69	
	HBA1_2	MN1_HBA1_2	MN1_HBA1_2_G900_1D	1D	70	
DS120 Server 2	HBA2_1	MN1_HBA1_1	MN1_HBA1_1_G900_2C	2C	69	
	HBA2_2	MN1_HBA1_2	MN1_HBA1_2_G900_2D	2D	70	

---

**Note** — In a production environment, it is recommended to use separate storage ports for the management servers to avoid impact on the database performance. Shared storage ports can be used, however port utilization should be monitored to avoid performance issues on high performance environments.

---

## Network Architecture

This architecture requires the following separate networks:

- **Private Network (also called cluster interconnect)** — This network must be scalable. In addition, it must meet the low latency needs of the network traffic generated by the cache synchronization of Oracle Real Application Clusters and inter-node communication among the nodes in the cluster.
- **Public Network** — This network provides client connections to the applications and Oracle Real Application Clusters.

Hitachi Vantara recommends using pairs of 25 Gbps NICs for the cluster interconnect network and public network.

Observe these points when configuring private and public networks in your environment:

- For each server in the clusterware configuration, use at least two identical, high-bandwidth, low-latency NICs for the interconnection.
- Use NIC bonding to provide failover and load balancing of interconnections within a server.
- Set all NICs to full duplex mode.
- Use at least two public NICs for client connections to the application and database.
- Use at least two private NICs for the cluster interconnection.

### Physical Network Configuration

Figure 3 shows the network configuration in this solution.

**Figure 3**

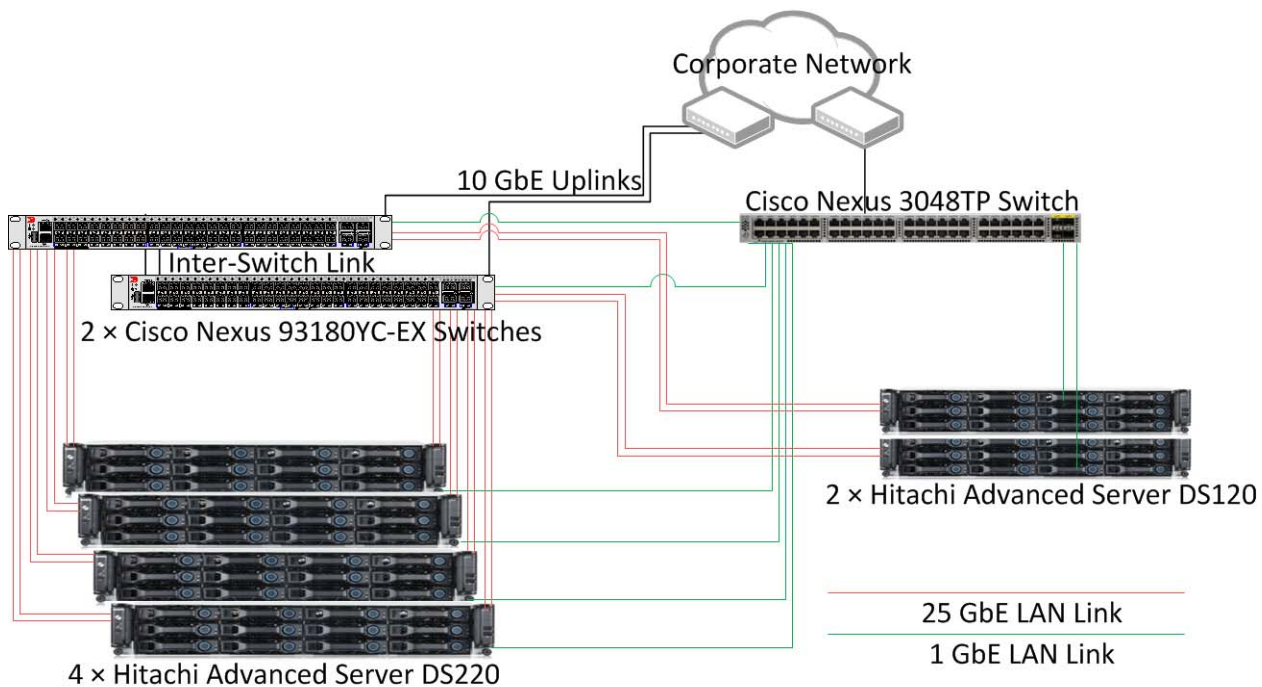


Table 11 shows the network configuration, and shows the virtual IP address and SCAN name configuration used when testing the environment. Your values may be different.

When creating NIC Bonding pairs, ports should be used on different cards to avoid single point of failure (SPoF).

TABLE 11. NETWORK CONFIGURATION

Server	NIC Ports	VLAN/ Subnet	NIC BOND	IP Address	Network	Bandwidth (Gbps)	Cisco Nexus 93180YC-EX Switch	
							Switch Number	Port
DS220 Server1	NIC - 0	208	Bond0	192.168.208.15	Private	25	1	41
	NIC - 2					25	2	
	NIC - 1	242	Bond1	192.168.242.23 3	Public Oracle	25	1	42
	NIC - 3					25	2	
	BMC- Dedicated NIC	244	-	192.168.244.16 1	Public Management	1	-	
DS220 Server2	NIC - 0	208	Bond0	192.168.208.16	Private	25	1	43
	NIC - 2					25	2	
	NIC - 1	242	Bond1	192.168.242.23 4	Public Oracle	25	1	44
	NIC - 3					25	2	
	BMC- Dedicated NIC	244	-	192.168.244.16 2	Public Management	1	-	
DS220 Server3	NIC - 0	208	Bond0	192.168.208.17	Private	25	1	45
	NIC - 2					25	2	
	NIC - 1	242	Bond1	192.168.242.23 5	Public Oracle	25	1	46
	NIC - 3					25	2	
	BMC- Dedicated NIC	244	-	192.168.244.16 3	Public Management	1	-	

TABLE 11. NETWORK CONFIGURATION (CONTINUED)

Server	NIC Ports	VLAN/ Subnet	NIC BOND	IP Address	Network	Bandwidth (Gbps)	Cisco Nexus 93180YC-EX Switch	
							Switch Number	Port
DS220 Server4	NIC - 0	208	Bond0	192.168.208.18	Private	25	1	47
	NIC - 2					25	2	
	NIC - 1	242	Bond1	192.168.242.236	Public Oracle	25	1	48
	NIC - 3					25	2	
	BMC-Dedicated NIC	244	-	192.168.244.164	Public Management	1	-	
DS120 management server1	NIC - 0	242	-	192.168.242.101	Public	25	1	49
	BMC-Dedicated NIC	244	-	192.168.244.169	Public Management	1	-	
DS120 management server2	NIC - 0	242	-	192.168.242.102	Public	25	1	50
	BMC-Dedicated NIC	244	-	192.168.244.170	Public Management	1	-	

TABLE 12. VIRTUAL IP AND SCAN NAME CONFIGURATION

Server	Virtual IP	Scan Name - hitachi-cluster-scan
Database Server 1 (DS220 1)	192.168.242.237	192.168.242.241 192.168.242.242 192.168.242.243
Database Server 1 (DS220 2)	192.168.242.238	
Database Server 3 (DS220 3)	192.168.242.239	
Database Server 4 (DS220 4)	192.168.242.240	



Table 13 lists virtual machine configuration running on management server cluster.

TABLE 13. MANAGEMENT SERVER VIRTUAL MACHINES CONFIGURATION

Virtual Machine	vCPU	Virtual Memory	Disk capacity	IP Address	OS
vCenter	2	10 GB	300 GB	192.168.242.102	VMware Photon Linux 1.0
OEM	16	32 GB	200 GB	192.168.242.16	RHEL 7.4
Oracle Adapters	2	6 GB	40 - 50 GB	192.168.242.80	OL 7.3
Oracle VM Manager	2	10 GB	100 GB	192.168.242.132	OL 7.3/7.4
HSA	4	16 GB	100 GB	192.168.242.81	CentOS 7.2
HIAA/HDCA	4	32 GB	800 GB	192.168.242.194	RHEL 7.3
HDCA Probe	4	10 GB	110 GB	192.168.242.197	RHEL 7.3
vSVP - Storage Virtual Platform G900	2	32	120 GB	192.168.167.49	Microsoft Windows 7 (64 bit)

### Hitachi Applications

The following are the Hitachi applications used for data analytics and performance monitoring during execution of this solution:

- Hitachi Infrastructure Analytics Advisor (HIAA)
- Hitachi Storage Advisor (HSA)

#### *Hitachi Infrastructure Analytics Advisor (HIAA)*

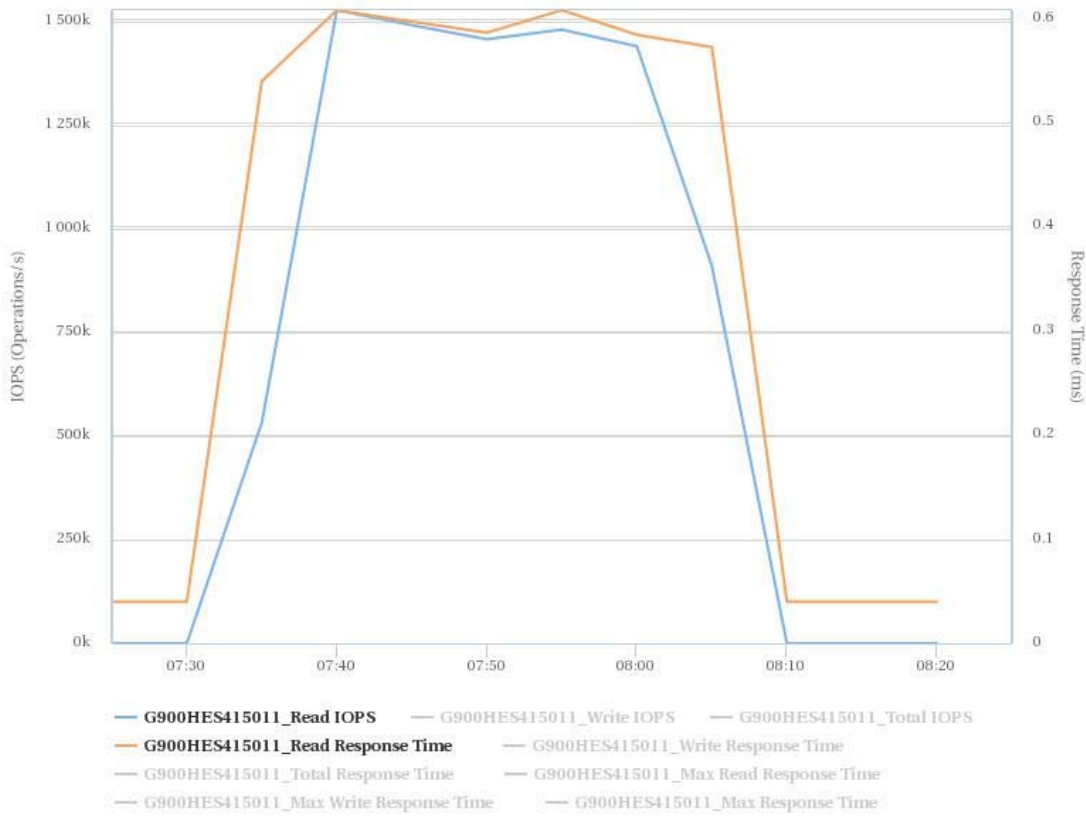
The following are the key features of Hitachi Infrastructure Analytics Advisor:

- Unified infrastructure monitoring dashboard
- Advanced reporting
- Storage I/O controls for SLO management
- System and Resource Events
- Granular Data Collection
- End-to-end monitoring

Please refer to the [Hitachi Infrastructure Analytics Advisor](#) for more details.

Figure 4 shows Hitachi Data Center Analytics performance page with **Random Read Orion test performance graph**.

Figure 4



### Hitachi Storage Advisor (HSA)

Hitachi Storage Advisor is a unified software management tool that reduces the complexity of managing storage systems by simplifying the setup, management, and maintenance of storage resources.

Some of the key Storage Advisor capabilities include:

- Simplified user experience for managing infrastructure resources.
- Recommended system configurations to speed initial storage system setup and accelerate new infrastructure resource deployments.
- Integrated configuration workflows with Hitachi Vantara recommended practices to streamline storage provisioning and data protection tasks.
- Common, centralized management for supported storage systems.
- A REST-based API to provide full management programmability and control in addition to unified file-based management support.
- Storage Advisor enables automated SAN zoning during volume attaches and detach. Optional auto-zoning eliminates the need for repetitive zoning tasks to be performed on the switch.

Please refer to the [Hitachi Storage Advisor](#) documentation for more details.

Figure 5 shows Hitachi Storage Advisor with VSP G900 storage system.

**Figure 5**

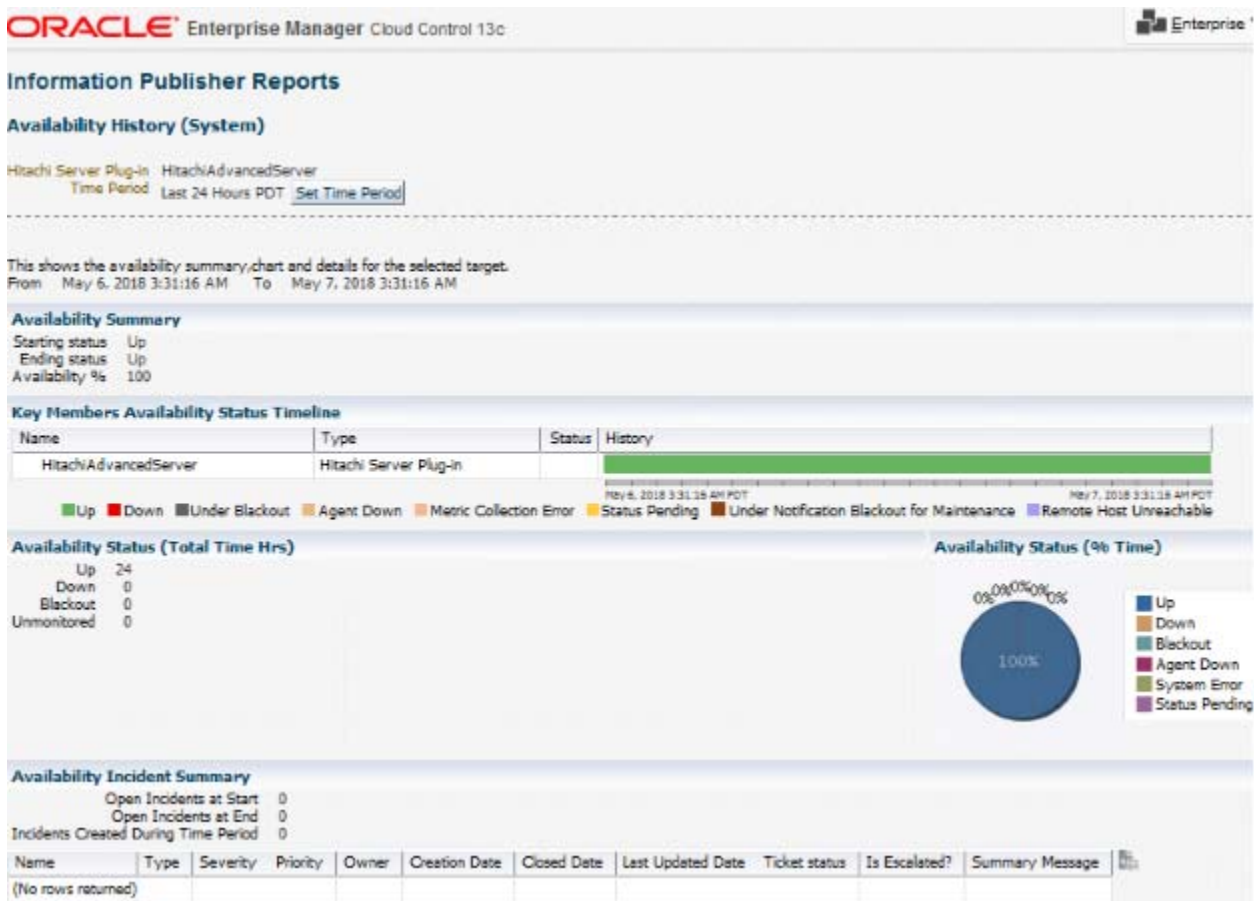


### *Hitachi Server Adapter for Oracle Enterprise Manager*

Hitachi Server Adapter for Oracle Enterprise Manager shows the availability summary chart and details for the selected target.

Figure 6 shows a System availability history report.

Figure 6



## Engineering Validation

This summarizes the key observations from the test results for the Hitachi Unified Compute Platform CI architecture for Oracle RAC deployment with Hitachi Virtual Storage Platform G900 and Hitachi Advanced Server DS220.

## Test Methodology

The test results are demonstrated using Oracle Orion and peakmarks tools.

### Oracle 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](#).

The Oracle Orion 12.1.0.2.0 tool is used to validate this solution. Orion tests are performed with straight storage mapping, which can yield better performance than traditional storage mapping. Straight storage mapping would limit the availability or redundancy in the architecture.

## Peakmarks

[peakmarks](#) is the leading benchmark software for Oracle platforms that is used 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 9.2 tool is used to validate this solution.

## Database Configuration

Table 14 shows parameter details for four-node Oracle Real Application Clusters ASM database.

TABLE 14. ORACLE DATABASE CONFIGURATION

Oracle Database Parameter	Value
compatible	12.1.0.2.0
cluster_database	TRUE
cluster_database_instances	4
Oracle Database size	8 TB
Database Storage Type	ASM
Database fill factor	80%

## Test Results

Table 15 lists the results of **Oracle Orion** test cases used to validate this solution.

TABLE 15. ORACLE ORION TEST RESULTS

Test Case	Test / Workload type	Metric	Value
1	Storage performance - 100% OLTP Random Read (8k)	Max. IO/s	1,448,845
		Avg. RT	0.98 ms
2	Storage performance - 100% OLTP Random Writes (8k)	Max. IO/s	312,350
		Avg. RT	1.15 ms

TABLE 15. ORACLE ORION TEST RESULTS (CONTINUED)

Test Case	Test / Workload type	Metric	Value
3	Storage performance - 100% OLAP Sequential Reads (1024K)	Max. Throughput	22.52 GB/sec
4	Storage performance - 100% OLAP Sequential Writes (1024K)	Max. Throughput	13.39 GB/sec

Table 16 lists the results of **peakmarks** test cases used to validate this solution.

TABLE 16. PEAKMARKS TEST RESULTS

Test Case	Test / Workload type	Metric	Value
1	Storage performance random read (STO-RR)	Max. IO/s	1,006,402
		Avg. RT	1.01 ms
2	Storage performance random write (STO-RWF)	Max. IO/s	378,041
		Avg. RT	0.44
3	Storage performance sequential read (STO-SR)	Max. Throughput	23.32 GB/s
4	Storage mixed random read write (STO-MIX 20% update ratio)	Max. IO/s	423,900
		Avg. RT	0.8 ms
5	Database medium OLTP select performance - 25 rows per transaction (DBX-S25)	Throughput in transactions per second	57,090
		Throughput in rows per second	1,427,275
		Avg. RT for SQL statement	1 ms

TABLE 16. PEAKMARKS TEST RESULTS (CONTINUED)

Test Case	Test / Workload type	Metric	Value
6	Server performance test - OLTP 25 rows per transaction(SRV-S25)	Throughput in transactions per second	16,878
		Throughput in rows per second	2,430,470
		Throughput in logical buffer reads per second	64,359,458
		Avg. RT for SQL statement	0.1 ms
7	CPU processor performance test - Arithmetic ADD operation (CP2-SA)	Throughput in operations per second	319,117,415,406

## For More Information

Hitachi Vantara Global Services offers experienced storage consultants, proven methodologies and a comprehensive services portfolio to assist you in implementing Hitachi products and solutions in your environment. For more information, see the [Services](#) website.

Demonstrations and other resources are available for many Hitachi products. To schedule a live demonstration, contact a sales representative or partner. To view on-line informational resources, see the [Resources](#) website.

Hitachi Academy is your education destination to acquire valuable knowledge and skills on Hitachi products and solutions. Our Hitachi Certified Professional program establishes your credibility and increases your value in the IT marketplace. For more information, see the Hitachi Vantara [Training and Certification](#) website.

For more information about Hitachi products and services, contact your sales representative, partner, or visit the [Hitachi Vantara](#) website.



## Hitachi Vantara



Corporate Headquarters  
2845 Lafayette Street  
Santa Clara, CA 96050-2639 USA  
[www.HitachiVantara.com](http://www.HitachiVantara.com) | [community.HitachiVantara.com](http://community.HitachiVantara.com)

Regional Contact Information  
**Americas:** +1 408 970 1000 or [info@hitachivantara.com](mailto:info@hitachivantara.com)  
**Europe, Middle East and Africa:** +44 (0) 1753 618000 or [info.emea@hitachivantara.com](mailto:info.emea@hitachivantara.com)  
**Asia Pacific:** +852 3189 7900 or [APACMarketing@hitachivantara.com](mailto:APACMarketing@hitachivantara.com)

© Hitachi Vantara Corporation 2018. All rights reserved. HITACHI is a trademark or registered trademark of Hitachi, Ltd. Microsoft and Windows are trademarks or registered trademarks of Microsoft Corporation. All other trademarks, service marks and company names are properties of their respective owners.

Notice: This document is for informational purposes only, and does not set forth any warranty, expressed or implied, concerning any equipment or service offered or to be offered by Hitachi Vantara.

MK-SL-062-01 May 2018.