

WHITE PAPER

# Hitachi Solution for Databases Reference Architecture for Oracle Real Application Clusters Database 12c

Reference Architecture Guide

By Amol Bhoite

November 2017

# 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              |
|-----------|-----------------|-------------------|
| SL-014-00 | Initial release | November 20, 2017 |

# 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 Gx00 Models              | 6         |
| Hitachi Advanced Server DS220 Server                      | 6         |
| Hitachi Advanced Server DS120 Server                      | 6         |
| Brocade Switches  | 6         |
| Cisco Switches  | 7         |
| Red Hat Enterprise Linux                                  | 7         |
| Oracle Database With the Real Application Clusters Option | 7         |
| Device Mapper Multipathing                                | 7         |
| Hitachi Infrastructure Analytics Advisor                  | 7         |
| Hitachi Storage Advisor                                   | 8         |
| Oracle Enterprise Manager                                 | 8         |
| Hitachi Storage Adapter for Oracle Enterprise Manager     | 8         |
| <b>Solution Design</b>                                    | <b>8</b>  |
| Storage Architecture                                      | 9         |
| Server and Application Architecture                       | 14        |
| SAN Architecture  | 15        |
| Network Architecture                                      | 16        |
| Hitachi Applications                                      | 18        |
| <b>Engineering Validation</b>                             | <b>24</b> |
| Test Methodology  | 24        |
| Test Results  | 26        |

# Hitachi Solution for Databases Reference Architecture for Oracle Real Application Clusters Database 12c

## 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 G800 (VSP G800), 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 on-line 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 7.3. 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 G800
- 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 two nodes using Hitachi Virtual Storage Platform G800. This environment addresses the high availability, performance, and scalability requirements for OLTP and OLAP workloads. Tailor your implementation of this solution 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.

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

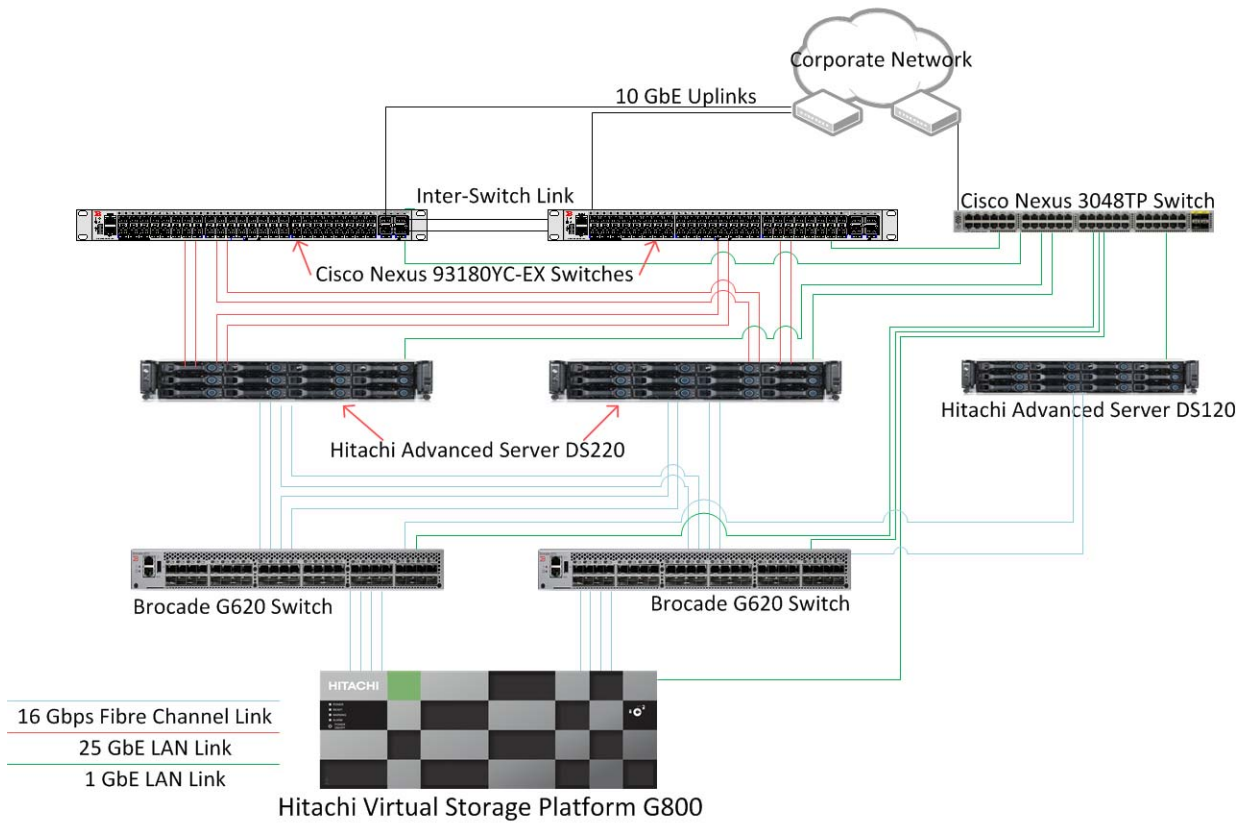
## High Level Infrastructure

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

The configuration of Virtual Storage Platform G800 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 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 and Table 2.

TABLE 1. HARDWARE COMPONENTS

| Hardware                              | Detailed Description   | Firmware/Version                        | Quantity |
|---------------------------------------|--|---|----------|
| Hitachi Virtual Storage Platform G800 | <ul style="list-style-type: none"> <li>■ Two Controllers</li> <li>■ 8 x 16 Gbps Fibre Channel Ports</li> <li>■ 8 x 12 Gbps Backend SAS Ports</li> <li>■ 512 GB cache memory</li> <li>■ 40 x 960 GB SSDs Plus 2 spare</li> <li>■ 16 Gbps x 2 ports CHB</li> </ul> | 83-04-47-60/00                          | 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 (64GB x 12) DIMM DDR4 Synchronous Registered (Buffered) 2666 MHz</li> </ul>  | 3A08.HX<br>X = latest supported version | 2        |
|                                       | <ul style="list-style-type: none"> <li>■ Dual Port 25 GbE NIC Intel cards</li> </ul>   | 5.51                                    | 2        |
|                                       | <ul style="list-style-type: none"> <li>■ Emulex LightPulse LPe31002-M6 2-Port 16 Gb Fibre Channel Adapter</li> </ul>   | FV11.2.156.27                           | 2        |
| Hitachi Advanced Server DS120 server  | <ul style="list-style-type: none"> <li>■ 2 x Intel Xeon Silver 4110 CPU @ 2.10GHz</li> <li>■ 128 GB (32 GB x 4) DIMM DDR4 Synchronous Registered (Buffered) 2666 MHz</li> <li>■ 2 x 128 GB MLC SATADOM for boot</li> </ul>                                       | 3A08.HX<br>X = latest supported version | 1        |
|                                       | <ul style="list-style-type: none"> <li>■ Dual Port 25 GbE NIC Intel cards</li> </ul>   | 5.51                                    | 1        |
|                                       | <ul style="list-style-type: none"> <li>■ Emulex LightPulse LPe31002-M6 2-Port 16 Gb Fibre Channel Adapters</li> </ul>  | FV11.2.156.27                           | 1        |

TABLE 1. HARDWARE COMPONENTS

| Hardware                       | Detailed Description   | Firmware/Version | Quantity |
|--------------------------------|--|------------------|----------|
| 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> | V8.0.1           | 2        |
| Cisco Nexus                    | <ul style="list-style-type: none"> <li>■ 93180YC-EX</li> <li>■ 48 × 10/25 GbE fiber ports</li> <li>■ 6 × 40/100 Gbps Quad SFP (QSFP28) ports</li> </ul>                      | ■ 7.0(3)I5(1)    | 2        |
| Cisco Nexus                    | <ul style="list-style-type: none"> <li>■ 3048TP</li> <li>■ 1 GE 48-Port Gb Ethernet Switch</li> </ul>  | ■ 7.0(3)I4(2)    | 2        |

TABLE 2. SOFTWARE COMPONENTS

| Software   | Version  | Function   |
|--|--|--|
| Red Hat Enterprise Linux                         | RHEL 7.3<br>(Kernel Version - 3.10.0-514.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  |
| Hitachi Storage Advisor (HSA)                    | 2.1.0  | Storage orchestration software   |
| Hitachi Infrastructure Analytics Advisor (HIAA)  | 2.1  | Analytics Software   |



TABLE 2. SOFTWARE COMPONENTS (CONTINUED)

| Software  | Version                  | Function  |
|---|--------------------------|---|
| Manager for Hitachi adapters for Oracle Database      | 2.0.0                    | Hitachi adapters management<br>Virtual appliance software |
| Hitachi Storage Adapter for Oracle Enterprise Manager | 2.0.0                    | Storage management software                               |
| Hitachi Server Adapter for Oracle Enterprise Manager  | 2.0.0                    | Storage 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                    |

### Hitachi Virtual Storage Platform Gx00 Models

[Hitachi Virtual Storage Platform Gx00 models](#) are based on industry-leading enterprise storage technology. With flash-optimized performance, these systems provide advanced capabilities previously available only in high-end storage arrays. With the Virtual Storage Platform Gx00 models, you can build a high performance, software-defined infrastructure to transform data into valuable information.

Hitachi Storage Virtualization Operating System provides storage virtualization, high availability, superior performance, and advanced data protection for all Virtual Storage Platform Gx00 models. This proven, mature software provides common features to consolidate assets, reclaim space, extend life, and reduce migration effort.

This solution uses Virtual Storage Platform G800, 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 unleashing computing performance, sophisticated power and thermal design to avoid unnecessary OPEX with quick deployment.

### 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 attempt to 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

## 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.

## 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.

## 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.

## 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 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.

## 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 G800.

The infrastructure configuration includes the following:

- **Oracle RAC Servers** - There are two 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.

## **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

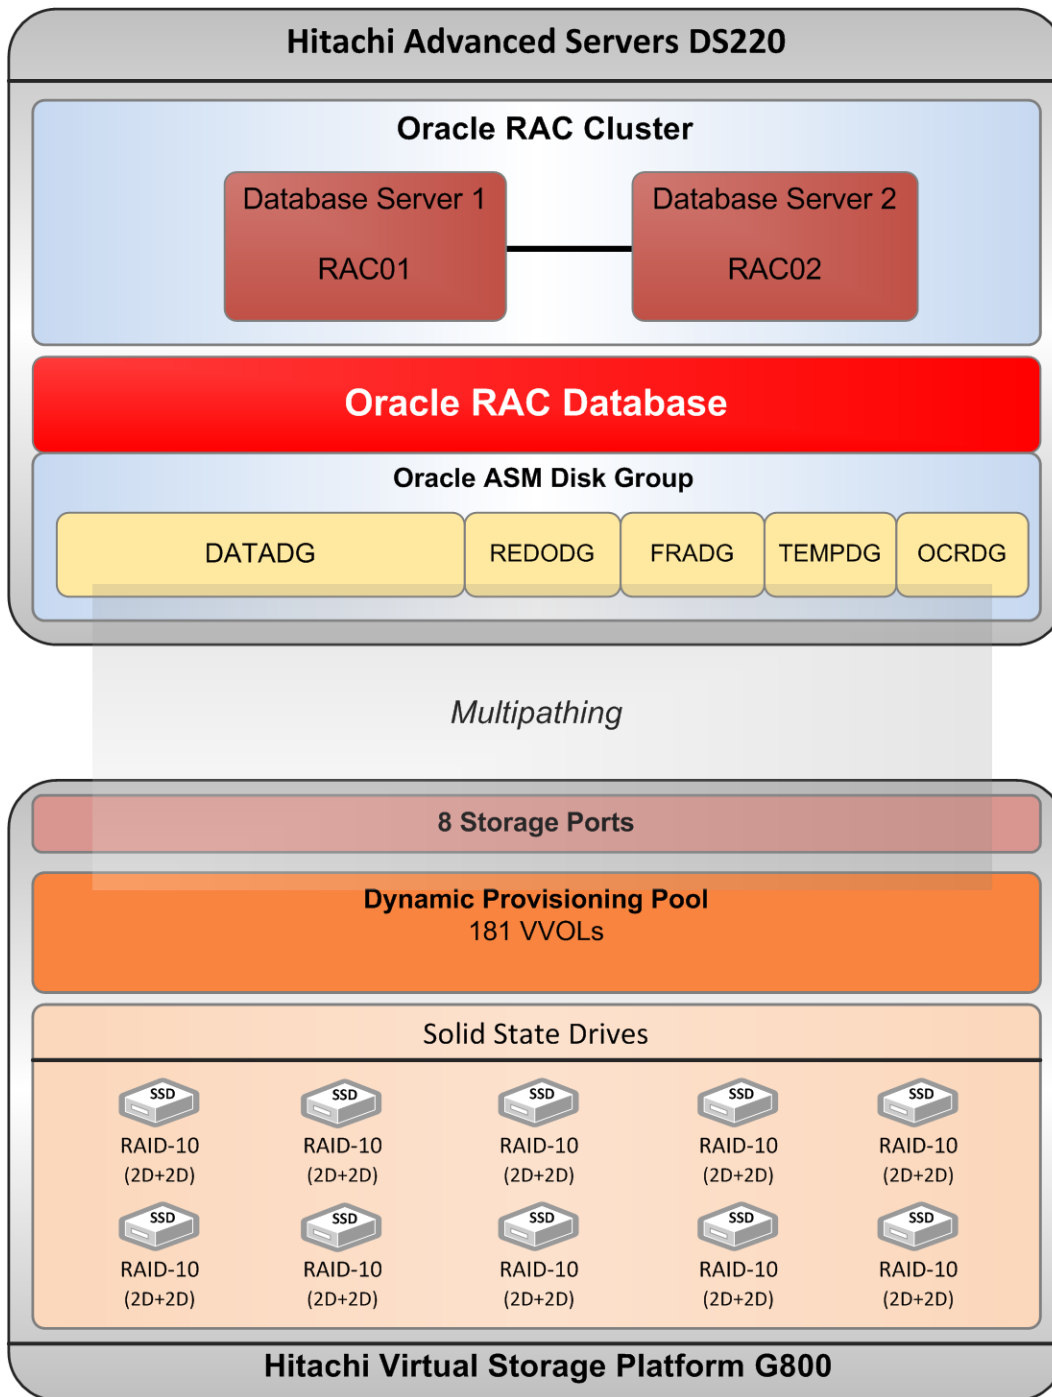


Table 3 shows the storage pool configuration used for this solution.

TABLE 3. STORAGE POOL CONFIGURATION

|                             |                                 |
|-----------------------------|---------------------------------|
| Pool ID                     | Unified Compute-Platform-Oracle |
| Pool Type                   | Dynamic Provisioning            |
| RAID Group                  | 1-1 – 1-10                      |
| RAID Level                  | RAID-10 (2D+2D)                 |
| Drive Type                  | 960 GB SSDs                     |
| Number of Drives            | 40                              |
| Number of Spare Drives      | 2                               |
| Number of Pool Volume LDEVs | 40                              |
| Pool Volume LDEV size       | 440 GB                          |
| Pool Capacity               | 17.18 TB                        |

Table 4 shows the logical storage configuration used in this solution.

TABLE 4. LOGICAL STORAGE CONFIGURATION

|                 |   |
|-----------------|---|
| Pool ID         | Unified Compute-Platform-Oracle   |
| Number of VVOLs | 181   |
| VVOL Size       | 2 × 200 GB, 128 × 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, 1B, 2B, 3A, 4A, 3B, 4B  |

### *Database Layout*

The database layout design uses recommended practices from Hitachi Vantara for Hitachi Virtual Storage Platform G800 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.

Table 5 shows the Oracle RAC Database Settings

TABLE 5. ORACLE RAC DATABASE SETTINGS

| <i>Environment</i> | <i>Value</i>              |
|--------------------|---------------------------|
| RAC configuration  | Yes                       |
| ASM                | Yes - Oracle RAC Database |

Table 6 shows the Oracle Environment Parameters.

TABLE 6. ORACLE ENVIRONMENT PARAMETERS

| <i>Setting</i>        | <i>Value</i> |
|-----------------------|--------------|
| 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_ASYNC_IO         | TRUE         |



Table 7 shows the details of the disk mappings from the LUNs to the ASM disk groups for Oracle RAC Database tablespaces.

TABLE 7. LUNs and Oracle ASM Disk Mappings

| ASM Disk Group | ASM Disk                   | DM-Multipath LUNs  | LUN Details     | Purpose                                 |
|----------------|----------------------------|--|-----------------|---|
| OCRDG          | OCRDISK1 -<br>OCRDISK3     | /dev/mapper/mpathaa - /dev/mapper/mpathac  | 3 x 5 GB        | Oracle Cluster Registry and Voting Disk |
| REDODG         | REDODISK1 -<br>REDODISK16  | /dev/mapper/mpathad - /dev/mapper/mpathap<br>/dev/mapper/mpathba - /dev/mapper/mpathbc   | 16 x 10 GB      | Online REDO log group                   |
| FRADG          | FRADISK1 -<br>FRADISK16    | /dev/mapper/mpathbd - /dev/mapper/mpathbp<br>/dev/mapper/mpathcb - /dev/mapper/mpathcc   | 16 x 40 GB      | Flash Recovery Area                     |
| TEMPDG         | TEMPDISK1 -<br>TEMPDISK16  | /dev/mapper/mpathcd - /dev/mapper/mpathcp<br>/dev/mapper/mpathda - /dev/mapper/mpathdc   | 16 x 40 GB      | Temp                                    |
| DATADG         | DATADISK1 -<br>DATADISK128 | /dev/mapper/mpathdd - /dev/mapper/mpathdp<br>/dev/mapper/mpatheb - /dev/mapper/mpathep<br>/dev/mapper/mpathfa - /dev/mapper/mpathfp<br>/dev/mapper/mpathga - /dev/mapper/mpathgp<br>/dev/mapper/mpathhd - /dev/mapper/mpathhp<br>/dev/mapper/mpathib - /dev/mapper/mpathip<br>/dev/mapper/mpathja - /dev/mapper/mpathjp<br>/dev/mapper/mpathka - /dev/mapper/mpathkp | 128 x<br>160 GB | Application Data                        |

## Server and Application Architecture

This reference architecture uses two Hitachi Advanced Server DS220 servers for a two-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 8 describes the details of the server configuration for this solution.

This reference architecture uses one Hitachi Advanced Server DS120 server for VMware ESXi management server configuration.

Details of the VMware ESXi management server are specified in Table 8.

TABLE 8. HITACHI ADVANCED SERVER DS220 AND DS120 SERVER SPECIFICATIONS

| Hitachi Advanced Server | Server            | Server Name   | Role  | CPU Core | RAM                 |
|-------------------------|-------------------|---------------|---|----------|---------------------|
| DS220                   | Oracle Server1    | oracle-rac-01 | Oracle RAC node 1                                   | 36       | 768 GB (64 GB × 12) |
|                         | Oracle Server2    | oracle-rac-02 | Oracle RAC node 2                                   | 36       | 768 GB (64 GB × 12) |
| DS120                   | Management server | VMware ESXi   | Hitachi Storage Advisor VM                          | 16       | 128 GB (32 GB × 4)  |
|                         |                   |               | Hitachi Infrastructure Analytics Advisor VM         |          |                     |
|                         |                   |               | Manager for Hitachi Adapters for Oracle Database VM |          |                     |
|                         |                   |               | Oracle Enterprise Manager Cloud Control 13c VM      |          |                     |

## SAN Architecture

Map the provisioned LDEVs to multiple ports on Hitachi Virtual Storage Platform G800 (VSP G800). These LDEV port assignments provide multiple paths to the storage system from the host for high availability.

- 8 SAN switch connections are being used for VSP G800 host ports.
- 8 SAN switch connections are being used for server HBA ports.

Table 9 shows details of the Fibre Channel switch connect configuration on the Hitachi Virtual Storage Platform G800 ports.

TABLE 9. SAN HBA CONNECTION CONFIGURATION TO VSP G800

| Server         | HBA Ports | Storage Host Group | Switch Zone                | Storage System | Storage Port | Brocade G620 Switch |
|----------------|-----------|--------------------|----------------------------|----------------|--------------|---------------------|
| DS220 Server 1 | HBA1      | HITACHI1_HBA1      | HITACHI1_HBA1_ASE42_183_1A | VSP G800       | 1A           | 69                  |
|                | HBA2      | HITACHI1_HBA2      | HITACHI1_HBA2_ASE42_183_2A |                | 2A           | 70                  |
|                | HBA3      | HITACHI1_HBA3      | HITACHI1_HBA3_ASE42_183_1B |                | 1B           | 69                  |
|                | HBA4      | HITACHI1_HBA4      | HITACHI1_HBA4_ASE42_183_2B |                | 2B           | 70                  |
| DS220 Server 2 | HBA1      | HITACHI2_HBA1      | HITACHI2_HBA1_ASE42_183_3A |                | 3A           | 69                  |
|                | HBA2      | HITACHI2_HBA2      | HITACHI2_HBA2_ASE42_183_4A |                | 4A           | 70                  |
|                | HBA3      | HITACHI2_HBA3      | HITACHI2_HBA3_ASE42_183_3B |                | 3B           | 69                  |
|                | HBA4      | HITACHI2_HBA4      | HITACHI2_HBA4_ASE42_183_4B |                | 4B           | 70                  |
| DS120 Server   | HBA1      | HDCA_HBA1          | HDCA_HBA1_ASE42_183_1A     |                | 1A           | 69                  |
|                | HBA2      | HDCA_HBA2          | HDCA_HBA2_ASE42_183_3A     |                | 3A           | 70                  |

## 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**

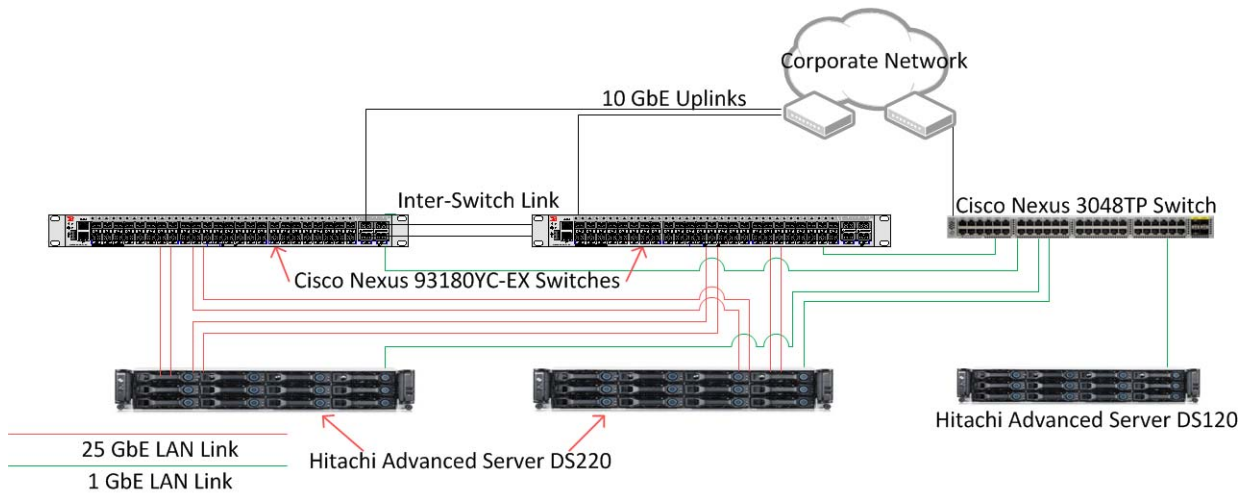


Table 10 shows the network configuration, and Table 11 shows the virtual IP address and SCAN name configuration used when testing the environment. Your values may be different. shows the network configuration, and Table 11 shows the virtual IP address and SCAN name configuration used when testing the 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 10. NETWORK CONFIGURATION**

| Server        | NIC Ports         | 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    | 172.17.242.233 | Public Oracle     | 25               | 1                             | 42   |
|               | NIC - 3           |        |          |                |                   | 25               | 2                             |      |
|               | BMC-Dedicated NIC | 244    | -        | 172.17.244.161 | Public Management | 1                | -                             | -    |

TABLE 10. NETWORK CONFIGURATION (CONTINUED)

|  |                   |     |       |                |                   |    |   |    |
|--|-------------------|-----|-------|----------------|-------------------|----|---|----|
| DS220<br>Server2                             | NIC - 0           | 208 | Bond0 | 192.168.208.16 | Private           | 25 | 1 | 43 |
|  | NIC - 2           |     |       |                |                   | 25 | 2 |    |
|  | NIC - 1           | 242 | Bond1 | 172.17.242.234 | Public Oracle     | 25 | 1 | 44 |
|  | NIC - 3           |     |       |                |                   | 25 | 2 |    |
|  | BMC-Dedicated NIC | 244 | -     | 172.17.244.162 | Public Management | 1  | - |    |
| DS120<br>Server3<br>manage<br>ment<br>server | BMC-Dedicated NIC | 244 | -     | 172.17.244.169 | Public Management | 1  | - |    |

TABLE 11. VIRTUAL IP AND SCAN NAME CONFIGURATION

| Server                      | Virtual IP     | Scan Name - hitachi-cluster-scan |
|-----------------------------|----------------|----------------------------------|
| Database Server 1 (DS220 1) | 172.17.242.236 | 172.17.242.240                   |
| Database Server 2 (DS220 2) | 172.17.242.237 | 172.17.242.241                   |
|                             |                | 172.17.242.242                   |

### 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 User Guide](#) for more details.

Figure 4 shows the Hitachi Infrastructure Analytics Advisor Dashboard tab for reference.

**Figure 4**

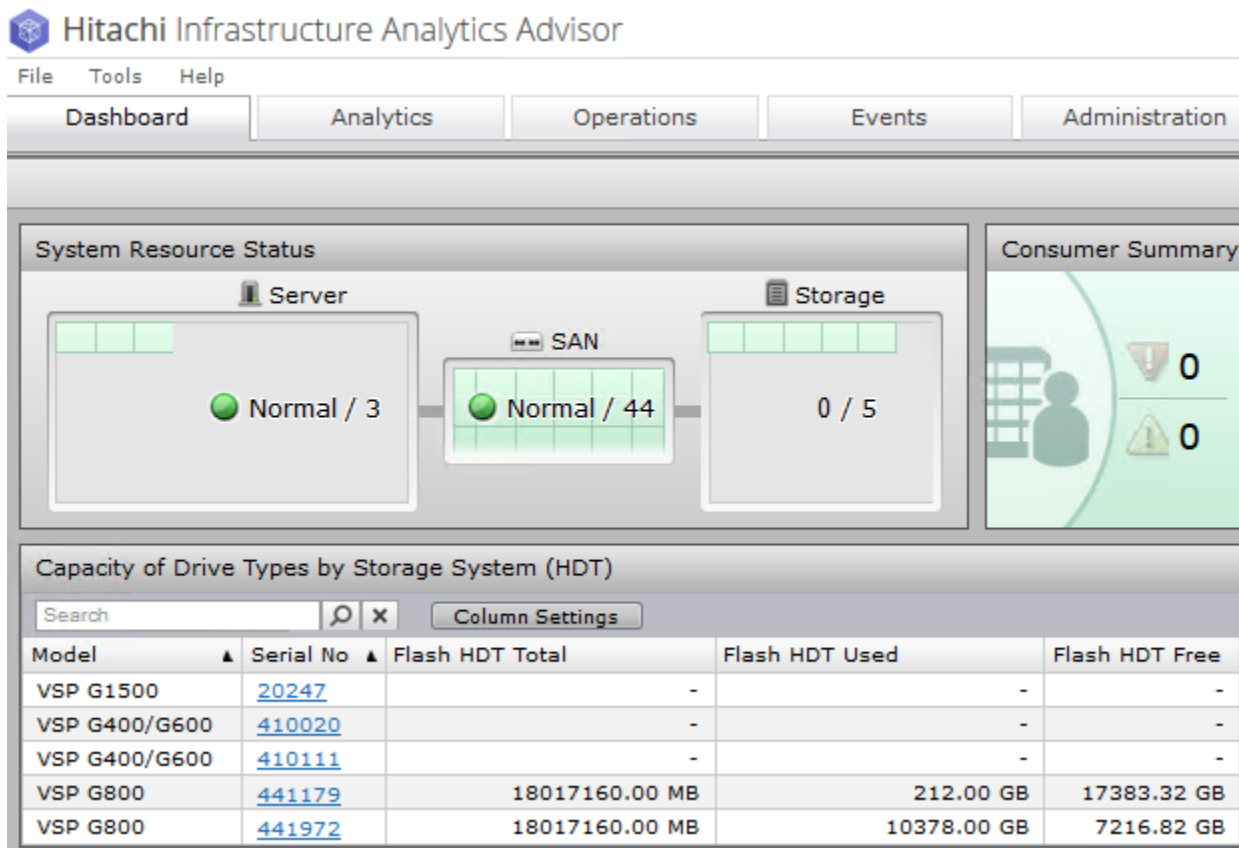
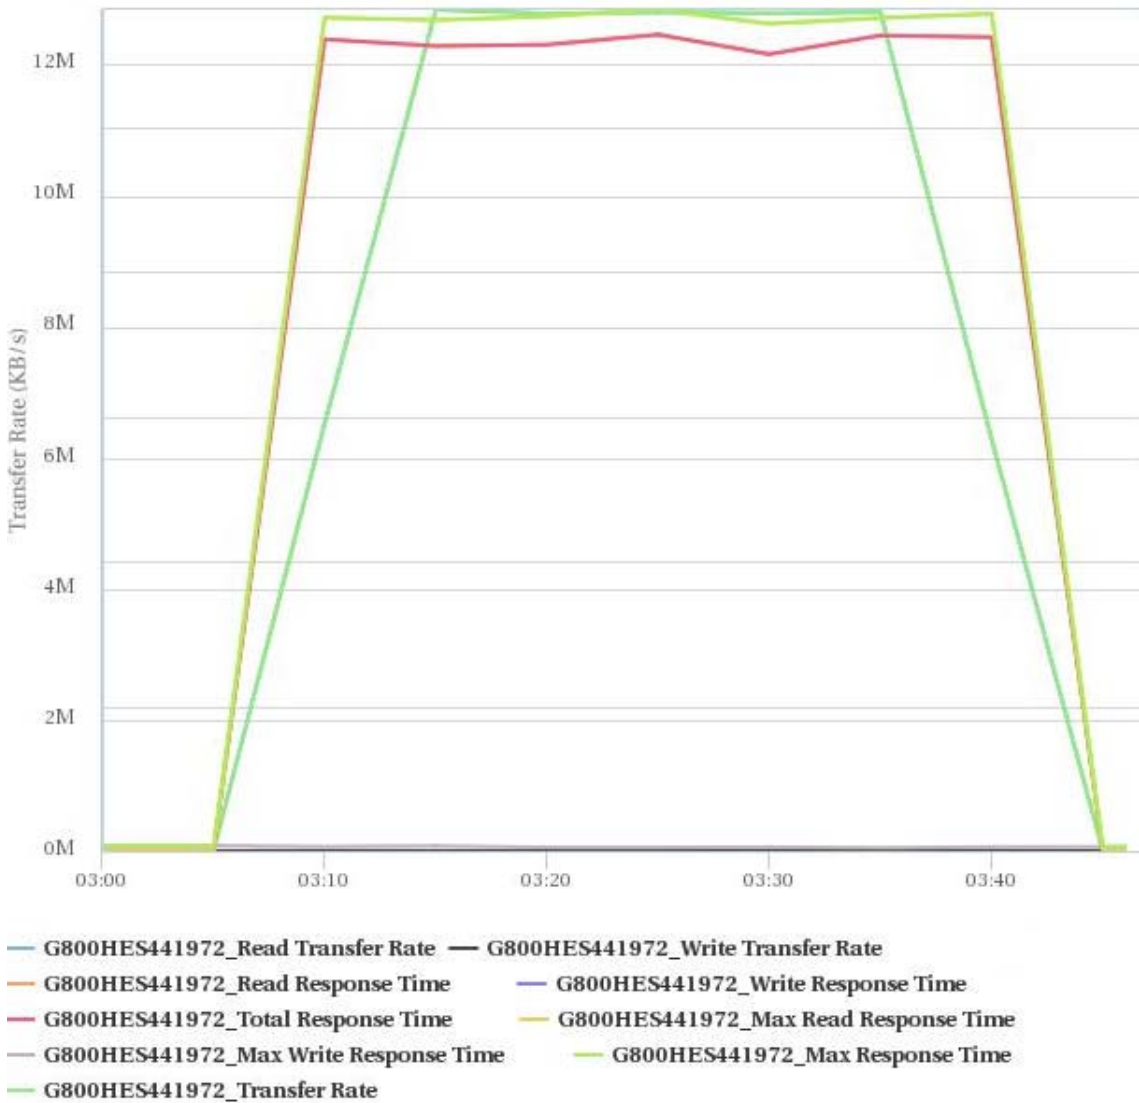


Figure 5 shows Hitachi Data Center Analytics performance page with Sequential Read Orion test performance graph.

**Figure 5**



### 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 User Guide](#) for more details.



Figure 6 shows Hitachi Storage Advisor with VSP G800 storage system.

**Figure 6**



### Hitachi Storage Adapter for Oracle Enterprise Manager

The Hitachi Storage Adapter for Oracle Enterprise Manager provides in-depth reports about different storage system components, mapping, configurations, usage, and performance summaries.

The adapter lets you monitor beyond the Oracle database object level to underlying Hitachi storage system objects. A comprehensive number of reports allows for a deeper understanding of your data and storage.

Figure 7 shows an ASM Disk Group to Storage Mapping report.

Figure 7

**ORACLE** Enterprise Manager Cloud Control 13c

### Information Publisher Reports

#### ASM Disk Group to Storage Mapping

Hitachi Storage Plug-in HitachiStorageAdaptorG800

This report provides the mapping between ASM disk groups and storage devices (LUNs). It also provides the storage utilization information for the ASM disk groups.

**Search**  
 Note: An empty filter lists all possible records.  
 Disk Group:   
 Storage Serial No.: All Storage   
 Host Name: All Hosts

**TIP** The search filter is case sensitive. Use '%' as a wildcard.

| Host Name    | ASM Disk Group | ASM Disk Group Size (GB) | Used Size (GB) | Free Size (GB) | ASM Disk    | ASM Disk Path         | Host Device | Storage Serial No. | LUN ID ( |
|--------------|----------------|--------------------------|----------------|----------------|-------------|-----------------------|-------------|--------------------|----------|
|              |                |                          |                |                | DATADG_0026 | /dev/mapper/mpathb1   | /dev/sdlv   |                    | 00:00:2D |
|              |                |                          |                |                | DATADG_0039 | /dev/mapper/mpathbm1  | /dev/sdlwg  |                    | 00:00:6C |
|              |                |                          |                |                | DATADG_0044 | /dev/mapper/mpathg1   | /dev/sdlx   |                    | 00:00:2F |
|              |                |                          |                |                | DATADG_0040 | /dev/mapper/mpathc1   | /dev/sdly   |                    | 00:00:30 |
|              |                |                          |                |                | DATADG_0041 | /dev/mapper/mpathd1   | /dev/sdlz   |                    | 00:00:31 |
|              |                |                          |                |                | DATADG_0046 | /dev/mapper/mpathi1   | /dev/sdus   |                    | 00:00:32 |
|              |                |                          |                |                | DATADG_0058 | /dev/mapper/mpathk1   | /dev/sdulb  |                    | 00:00:33 |
|              |                |                          |                |                | DATADG_0043 | /dev/mapper/mpathf1   | /dev/sduc   |                    | 00:00:34 |
|              |                |                          |                |                | DATADG_0045 | /dev/mapper/mpathh1   | /dev/sduf   |                    | 00:00:35 |
|              |                |                          |                |                | DATADG_0112 | /dev/mapper/mpathm1   | /dev/sdus   |                    | 00:00:36 |
|              |                |                          |                |                | DATADG_0117 | /dev/mapper/mpathp1   | /dev/sduf   |                    | 00:00:37 |
|              |                |                          |                |                | DATADG_0047 | /dev/mapper/mpathj1   | /dev/sdug   |                    | 00:00:38 |
|              |                |                          |                |                | DATADG_0085 | /dev/mapper/mpathl1   | /dev/sduh   |                    | 00:00:39 |
|              |                |                          |                |                | DATADG_0118 | /dev/mapper/mpathq1   | /dev/sdui   |                    | 00:00:3A |
|              |                |                          |                |                | DATADG_0120 | /dev/mapper/mpaths1   | /dev/sduj   |                    | 00:00:3B |
|              |                |                          |                |                | DATADG_0115 | /dev/mapper/mpathn1   | /dev/sduk   |                    | 00:00:3C |
|              |                |                          |                |                | DATADG_0116 | /dev/mapper/mpatho1   | /dev/sdul   |                    | 00:00:3D |
|              |                |                          |                |                | DATADG_0124 | /dev/mapper/mpathw1   | /dev/sdum   |                    | 00:00:3E |
|              |                |                          |                |                | DATADG_0123 | /dev/mapper/mpathv1   | /dev/sdun   |                    | 00:00:3F |
|              |                |                          |                |                | DATADG_0119 | /dev/mapper/mpathr1   | /dev/sduo   |                    | 00:00:40 |
|              |                |                          |                |                | DATADG_0121 | /dev/mapper/mpatht1   | /dev/sdup   |                    | 00:00:41 |
|              |                |                          |                |                | DATADG_0001 | /dev/mapper/mpathel1  | /dev/sduq   |                    | 00:00:42 |
|              |                |                          |                |                | DATADG_0127 | /dev/mapper/mpathz1   | /dev/sdur   |                    | 00:00:43 |
|              |                |                          |                |                | DATADG_0122 | /dev/mapper/mpathu1   | /dev/sdus   |                    | 00:00:44 |
|              |                |                          |                |                | DATADG_0125 | /dev/mapper/mpathoc1  | /dev/sdut   |                    | 00:00:45 |
|              |                |                          |                |                | DATADG_0006 | /dev/mapper/mpathg1   | /dev/sduz   |                    | 00:00:4B |
| rac01-oracle | DATADG         | 10239.88                 | 2.52           | 10237.36       | DATADG_0000 | /dev/mapper/mpatheas1 | /dev/sdva   | 441972             | 00:00:4C |
|              |                |                          |                |                | DATADG_0007 | /dev/mapper/mpathel1  | /dev/sdvb   |                    | 00:00:4D |
|              |                |                          |                |                | DATADG_0008 | /dev/mapper/mpathel1  | /dev/sdvc   |                    | 00:00:4E |
|              |                |                          |                |                | DATADG_0011 | /dev/mapper/mpathel1  | /dev/sdvd   |                    | 00:00:4F |

### Hitachi Server Adapter for Oracle Enterprise Manager

Hitachi Server Adapter for Oracle Enterprise Manager shows information for all Hitachi Advanced servers that you configured for the Hitachi Server Adapter target.

Figure 8 shows Hitachi Advanced Server Report.

**Figure 8**

**ORACLE** Enterprise Manager Cloud Control 13c

### Information Publisher Reports

This report provides the information for all Quanta Servers that are configured for this Hitachi Server Adapter target.

---

#### Quanta Server Dashboard Information

Search Filter  
 Leave the filter set to 'All Servers' to see information for all Servers.  
 Server IP Address:

**TIP** Choose a Server from the pull-down filter menu to limit the report to that Server.

| Server IP Address | Firmware Revision             | Firmware Build Time      | BIOS Version | BMC NIC       | BMC Chipset | Host Name       | MAC Address       | IPV4 Address   |
|-------------------|-------------------------------|--------------------------|--------------|---------------|-------------|-----------------|-------------------|----------------|
| 172.17.242.161    | BMC Version : 3.16, IPMI v2.0 | Aug 28 2017 17:18:28 CST | 3A03.E5      | Dedicated NIC | AST2500(A2) | QCTA81E84968900 | A8:1E:84:96:89:00 | 172.17.242.161 |

---

#### Quanta Server Performance Information

Search Filter  
 Leave the filter set to 'All Servers' to see information for all Servers.  
 Server IP Address:

**TIP** Choose a Server from the pull-down filter menu to limit the report to that Server.

| Server IP Address | BMC Memory in Use (KB) | BMC Memory in Use (%) | BMC Memory Size (KB) | BMC Current Num Processes | BMC Current Num Users | BMC Processor Load (%) |
|-------------------|------------------------|-----------------------|----------------------|---------------------------|-----------------------|------------------------|
| 172.17.242.161    | 104652                 | 26                    | 394328               | 70                        | 0                     |                        |

---

#### Quanta Server Sensor Information

Search Filter  
 Leave the filter set to 'All Servers' to see information for all Servers.  
 Server IP Address:

**TIP** Choose a Server from the pull-down filter menu to limit the report to that Server.

Previous | 1-10 of 59 | ..

| Server IP Address | Sensor Name     | Sensor Type    | Status         | Current Reading | Lower Critical (LC) | Lower Non-Critical (LNC) | Upper Non-Critical | Upper Critical (UC) |
|-------------------|-----------------|----------------|----------------|-----------------|---------------------|--------------------------|--------------------|---------------------|
|                   | Fan_SYS0_0      | 172.17.242.161 | OK             | 2700 RPM        | 500.0               | 1000.0                   | 0.0                | 0.0                 |
|                   | Fan_SYS0_1      | 172.17.242.161 | OK             | 2500 RPM        | 500.0               | 1000.0                   | 0.0                | 0.0                 |
|                   | Fan_SYS0_Status | 172.17.242.161 | Device Present | -1 RPM          | 0.0                 | 0.0                      | Inf                | Inf                 |
|                   | Fan_SYS1_0      | 172.17.242.161 | OK             | 2700 RPM        | 500.0               | 1000.0                   | 0.0                | 0.0                 |
|                   | Fan_SYS1_1      | 172.17.242.161 | OK             | 2500 RPM        | 500.0               | 1000.0                   | 0.0                | 0.0                 |
|                   | Fan_SYS1_Status | 172.17.242.161 | Device Present | -1 RPM          | 0.0                 | 0.0                      | Inf                | Inf                 |
|                   | Fan_SYS2_0      | 172.17.242.161 | OK             | 2700 RPM        | 500.0               | 1000.0                   | 0.0                | 0.0                 |
|                   | Fan_SYS2_1      | 172.17.242.161 | OK             | 2500 RPM        | 500.0               | 1000.0                   | 0.0                | 0.0                 |
|                   | Fan_SYS2_Status | 172.17.242.161 | Device Present | -1 RPM          | 0.0                 | 0.0                      | Inf                | Inf                 |
|                   | Fan_SYS3_0      | 172.17.242.161 | OK             | 2800 RPM        | 500.0               | 1000.0                   | 0.0                | 0.0                 |

## 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 G800 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 12 shows parameter details for two-node Oracle Real Application Clusters ASM database.

TABLE 12. ORACLE DATABASE CONFIGURATION

| Oracle Database Parameter  | Value      |
|----------------------------|------------|
| compatible                 | 12.1.0.2.0 |
| cluster_database           | TRUE       |
| cluster_database_instances | 2          |
| Oracle Database size       | 8 TB       |
| Database Storage Type      | ASM        |
| Database fill factor       | 80%        |

## Test Results

Table 13 lists the results of **Oracle Orion** test cases used to validate this solution.

TABLE 13. ORACLE ORION TEST RESULTS

| <i>Test Case</i> | <i>Test / Workload type</i>                                  | <i>Metric</i>   | <i>Value</i> |
|------------------|--|-----------------|--------------|
| 1                | Storage performance -<br>100% OLTP Random Read (8k)          | Max. IO/s       | 601,499      |
|                  |  | Avg. RT         | 1.07 ms      |
| 2                | Storage performance -<br>100% OLTP Random Writes (8k)        | Max. IO/s       | 159,307      |
|                  |  | Avg. RT         | 1.1 ms       |
| 3                | Storage performance -<br>100% OLAP Sequential Reads (1024K)  | Max. Throughput | 12.21 GB/sec |
| 4                | Storage performance -<br>100% OLAP Sequential Writes (1024K) | Max. Throughput | 7.3 GB/sec   |

Table 14 lists the results of **Peakmarks** test cases used to validate this solution.

TABLE 14. PEAKMARKS TEST RESULTS

| <i>Test Case</i> | <i>Test / Workload type</i>                                | <i>Metric</i>   | <i>Value</i> |
|------------------|--|-----------------|--------------|
| 1                | Storage performance random read (STO-RR)                   | Max. IO/s       | 400,041      |
|                  |  | Avg. RT         | 1.07 ms      |
| 2                | Storage performance random write (STO-RWF)                 | Max. IO/s       | 200,094      |
|                  |  | Avg. RT         | 1.02 ms      |
| 3                | Storage performance sequential read (STO-SR)               | Max. Throughput | 10.58 GB/sec |
| 4                | Storage mixed random read write (STO-MIX 20% update ratio) | Max. IO/s       | 261,134      |
|                  |  | Avg. RT         | 1 ms         |



TABLE 14. PEAKMARKS TEST RESULTS (CONTINUED)

|   |   |   |                        |
|---|---|---|------------------------|
| 5 | Database medium OLTP select performance - 25 rows per transaction (DBX-S25) | Throughput in transactions per second         | 25,388                 |
|   |   | Throughput in rows per second                 | 634,694                |
|   |   | Avg. RT for SQL statement                     | 1.1 ms                 |
| 6 | Server performance test - OLTP 25 rows per transaction(SRV-S25)             | Throughput in transactions per second         | 17,950                 |
|   |   | Throughput in rows per second                 | 1,292,394              |
|   |   | Throughput in logical buffer reads per second | 33,731,989             |
|   |   | Avg. RT for SQL statement                     | 0.1                    |
| 7 | CPU processor performance test - Arithmetic ADD operation (CP2-SA)          | Throughput in operations per second           | <b>165,090,903,205</b> |

## 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 [hds.marketing.apac@hitachivantara.com](mailto:hds.marketing.apac@hitachivantara.com)

© Hitachi Vantara Corporation 2017. All rights reserved. HITACHI is a trademark or registered trademark of Hitachi, Ltd. All other trademarks, service marks and company names are properties of their respective owners.

Hitachi Vantara Confidential: Internal use only.

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.

SL-014-00, November 2017.