

WHITE PAPER

# Hitachi Unified Compute Platform 6000 for Oracle Real Application Clusters using Hitachi Virtual Storage Platform G600 with Intel Xeon 2699v4 Processors

## Reference Architecture Guide

By Solutions Engineering (authors on next page)

July 2017

# Feedback

Hitachi Data Systems welcomes your feedback. Please share your thoughts by sending an email message to [SolutionLab@hds.com](mailto:SolutionLab@hds.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
AS-592-00	Initial release	April 3, 2017
AS-592-01	Update the authors	May 10, 2017
AS-592-02	Clarified information about disk drives.	July 3, 2017

# Authors

This document was authored by the following people:

- Libo Jiang
- Younus Choudhary
- Yoshihito Katayama
- Yuya Ohta
- Kazuya Sakamoto

# Table of Contents

- Solution Overview ..... 2**
- Key Solution Components ..... 3**
  - Hitachi Compute Blade 2500..... 4
  - Hitachi Virtual Storage Platform Gx00 Models ..... 4
  - Hitachi Accelerated Flash ..... 5
  - Hitachi Storage Virtualization Operating System..... 5
  - Hitachi Dynamic Link Manager Advanced ..... 6
  - Brocade Networking ..... 6
  - Red Hat Enterprise Linux ..... 6
  - Oracle Database ..... 6
- Solution Design ..... 7**
  - Storage Architecture..... 7
  - Database Layout..... 10
  - Server and Application Architecture ..... 12
  - SAN Architecture ..... 13
  - Network Architecture ..... 15
- Engineering Validation ..... 17**
  - Test Methodology ..... 17
  - Test Results ..... 18

# Hitachi Unified Compute Platform 6000 for Oracle Real Application Clusters using Hitachi Virtual Storage Platform G600 with Intel Xeon 2699v4 Processors

## Reference Architecture Guide

Use this reference architecture guide to see how Hitachi Unified Compute Platform 6000 for Oracle Real Application Cluster (UCP for Oracle RAC) provides a high performance, integrated converged solution for Oracle. The environment uses Hitachi Virtual Storage Platform G600 (VSP G600) with Intel Xeon 2699v4 processors (Broadwell) and Hitachi Accelerated Flash (HAF). With this, design an Oracle converged infrastructure to meet your requirements and budget.

This proven solution to optimize your Oracle database environment integrates servers, storage systems, network, and storage software. 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 server runs Oracle Database 12c R1 with the Oracle Real Application Cluster (RAC) option. The operating system is Red Hat Enterprise Linux 7.2.

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 G600 using Hitachi Accelerated Flash
- Hitachi Compute Blade 2500 (CB 2500)
- Storage area networks
- Oracle RAC Database 12c Release 1
- Oracle Automatic Storage Management (Oracle ASM)
- Red Hat Enterprise Linux
- Hitachi Dynamic Link Manager v8.4 (HDLM)

---

**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 6000 for Oracle Real Application Cluster on two nodes using Hitachi Virtual Storage Platform G600 with Hitachi Accelerated Flash. 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.

This reference architecture includes the following:


- **Hitachi Compute Blade 2500** with two 520H B4 server blades
  - **Server Blade 1** — Oracle RAC NODE 1
  - **Server Blade 2** — Oracle RAC NODE 2
- **Hitachi Virtual Storage Platform G600** with Hitachi Accelerated Flash
- 16 Gb/sec switched SAN infrastructure
- 10 GbE LAN infrastructure

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

The configuration of Virtual Storage Platform G600 and Compute Blade 2500 has the following:

- Fully redundant hardware
- Dual Fabric connectivity between hosts and storage
- Maintains high performance

**Figure 1**



To avoid any performance impact to the production database, Hitachi Data Systems 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

## Key Solution Components

Table 1 list the key hardware components used by this reference architecture.

**TABLE 1. KEY HARDWARE COMPONENTS**

Hardware	Detail Description	Version	Quantity
Hitachi Virtual Storage Platform G600	Dual controller	83-03-23-40/00	1
	16 × 16 Gb/sec Fibre Channel ports		
	16 backend serial attached SCSI (SAS) ports		
	256 GB cache memory		
	1.6 TB Flash memory drives (FMDs)		16
	1 spare		
	1.2 TB 10k RPM SAS Drives		24
Hitachi Compute Blade 2500 chassis	2 × 10 Gb/s DCB LAN switch module 10 Fan modules 2 Management modules 4 Power supply modules	Management Module Firmware Version A0122-B-1009	1
520H B4 Half Server Blade	2 Intel Xeon E5-2699v4 CPU	10-04/10-04	2
	256 GB (16 GB × 16) DDR4		
	1 × 4-port 10 GbE converged network adapter (CNA) LAN on motherboard (LOM)		
	Hitachi 2-port 16Gb/sec Fibre Channel Adapter	40-05-00	4
Brocade 6510 switch	48 port Fibre Channel switch	7.4.1b	1*
	16 Gb/s SFPs		

\*Although a single Brocade SAN switch was used when testing this environment, Hitachi Data Systems recommends using two switches for high availability purposes.

Table 2 list the key software components used in this reference architecture.

**TABLE 2. KEY SOFTWARE COMPONENTS**

Software	Version	Function
Red Hat Enterprise Linux	7.2	Operating system for Oracle RAC
Oracle ASM	12c Release 1 (12.1.0.2.0)	Volume manager and file system
Oracle Database	12c Release 1 (12.1.0.2.0)	Database system with multitenant architecture
Hitachi Storage Virtualization Operating System (SVOS) <ul style="list-style-type: none"> <li>■ Hitachi Dynamic Tiering</li> <li>■ Hitachi Dynamic Provisioning</li> </ul>	6.4.1	Storage virtualization, thin provisioning, storage service level controls, dynamic provisioning, and performance instrumentation
Hitachi Dynamic Link Manager	8.4.1	Multipath software

## Hitachi Compute Blade 2500

[Hitachi Compute Blade 2500](#) delivers enterprise computing power and performance with unprecedented scalability and configuration flexibility. Lower your costs and protect your investment.

Flexible I/O architecture and logical partitioning allow configurations to match application needs exactly with Hitachi Compute Blade 2500. Multiple applications easily and securely co-exist in the same chassis.

Add server management and system monitoring at no cost with Hitachi Compute Systems Manager. Seamlessly integrate with Hitachi Command Suite in Hitachi storage environments.

In this solution, Hitachi Compute Blade 2500 with 520H B4 server blades provides scalability and flexibility for an Oracle RAC configuration while providing redundancy.

## 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.

Virtual Storage Platform G600, used in this reference architecture, supports [Oracle Real Application Clusters](#).

## Hitachi Accelerated Flash

[Hitachi Accelerated Flash](#) features a flash module built specifically for enterprise-class workloads. Developed for Hitachi Virtual Storage Platform, Accelerated Flash is available for Hitachi Unified Storage VM and Hitachi Virtual Storage Platform family.

Accelerated Flash features innovative Hitachi-developed embedded flash memory controller technology. Hitachi flash acceleration software speeds I/O processing to increase flash device throughput.

Hitachi Accelerated Flash provides a reliable data storage for the Oracle database file placement with fast data retrieval for the OLTP workload.

## Hitachi Storage Virtualization Operating System

[Hitachi Storage Virtualization Operating System](#) (SVOS) spans and integrates multiple platforms. It integrates storage system software to provide system element management and advanced storage system functions. Used across multiple platforms, Storage Virtualization Operating System includes storage virtualization, thin provisioning, storage service level controls, dynamic provisioning, and performance instrumentation.

Storage Virtualization Operating System includes standards-based management software on a Hitachi Command Suite (HCS) base. This provides storage configuration and control capabilities for you.

Storage Virtualization Operating System uses Hitachi Dynamic Provisioning (HDP) to provide wide striping and thin provisioning. Dynamic Provisioning provides one or more wide-striping pools across many RAID groups. Each pool has one or more dynamic provisioning virtual volumes (DP-VOLs) without initially allocating any physical space. Deploying Dynamic Provisioning avoids the routine issue of hot spots that occur on logical devices (LDEVs).

This solution uses multiple dedicated dynamic provisioning pools for different types of Oracle data to avoid intermixing of the different types of data I/O. This benefits performance during peak database operations.

This solution uses [Hitachi Dynamic Tiering](#), a part of Storage Virtualization Operating System, for data mobility. Separately licensed, Dynamic Tiering virtualizes and automates mobility between tiers for maximum performance and efficiency. Instead of manually provisioning space from several storage technologies with different performance and cost characteristics, Dynamic Tiering manages multiple storage tiers as a single entity. It presents a virtual volume with embedded smart tiering to monitor access and move data based on demand.

Dynamic Tiering automatically moves infrequently referenced data to lower cost tiers of storage. This data placement provides higher performance with lower costs. It also provides automatic wide-striping performance optimization.

When Tier 1 hardware uses a solid state device or a flash module drive, active flash in Dynamic Tiering provides special care for write endurance. This complements using [flash storage](#) in your environment.

The dynamic tiering pool used for this solution provides the capability to later introduce different types of disks to increase the disk pool capacity that is used for the Oracle data without sacrifice of performance. Frequently accessed data is stored on flash modules.



## Hitachi Dynamic Link Manager Advanced

[Hitachi Dynamic Link Manager Advanced](#) combines all the capabilities of Hitachi Dynamic Link Manager and Hitachi Global Link Manager into a comprehensive multipathing solution. It includes capabilities such as the following:

- Path failover and failback
- Automatic load balancing to provide higher data availability and accessibility

This solution uses the **extended round-robin load balancing policy** configuration in Hitachi Dynamic Link Manager Advanced. This policy determines which path to allocate based on whether the data of the I/O to be issued is sequential with the immediately preceding data issued of the I/O.

If the data is sequential with the preceding data, the extended round-robin load balancing policy normally issues the same path for distributing the I/O that was issued to the immediately preceding I/O. However, after issuing a specified number of I/Os to a path, this policy then switches processing to the next path.

If the data is not sequential with the preceding data, the extended round-robin load balancing policy issues a path for distributing the I/O in order from among all the connected paths. Balancing the load across all available paths optimizes IOPS and response time.

## Brocade Networking

Brocade and Hitachi Data Systems 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.

The solution using the following Brocade products:

- Brocade VDX 2746 10 GbE DCB switch module
- Brocade 6510 48 port 16 Gb/s 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

[Oracle Database](#) has a multitenant 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.

## Solution Design

This describes the reference architecture environment to implement a quarter-rack environment for Hitachi Unified Compute Platform 6000 for Oracle RAC. The environment uses Hitachi Virtual Storage Platform G600 with Hitachi Accelerated Flash.

The infrastructure configuration includes the following:

- **Server** — There are two server nodes configured in an Oracle Real Application Cluster.
- **Storage System** — There are V-VOLs mapped to each port that are presented to the server as LUNs.
- **SAN Connection** — There is a SAN connection to connect the Fiber Channel HBA ports to the Brocade 6510 switch.

## Storage Architecture

This describes the storage architecture of this reference architecture. It takes into consideration Hitachi Data Systems and Oracle recommended practices for the deployment of database storage design.

### *Storage Layout*

The storage layout design uses recommended practice from Hitachi Data Systems for Hitachi Virtual Storage Platform G600 using Hitachi Accelerated Flash for better IOPS and throughput for OLTP operations.

- **Data Disks** — This structure has two different set of disks:
  - 1.6 TB flash module drive (FMD)
  - 1.2 TB 10k RPM SAS drives
- **Disk Configuration** — Create two types of parity groups using FMD and SAS drives. Assign the following RAID levels:
  - **FMD drives** — Four RAID-10 (2D+2D) parity groups
  - **SAS drives** — Four RAID-6 (6D+2P) parity groups
- **FMD Disk Tray Layout** — Use four disk trays. Use four drives from each disk tray. For each parity group, use one drive from each tray. This improves IOPS and for throughput.

## Storage Configuration

This is the high-level storage configuration diagram of this solution.

Figure 2 describes the storage configuration for this solution. This environment uses Hitachi Dynamic Link Manager for multipathing.

Figure 2

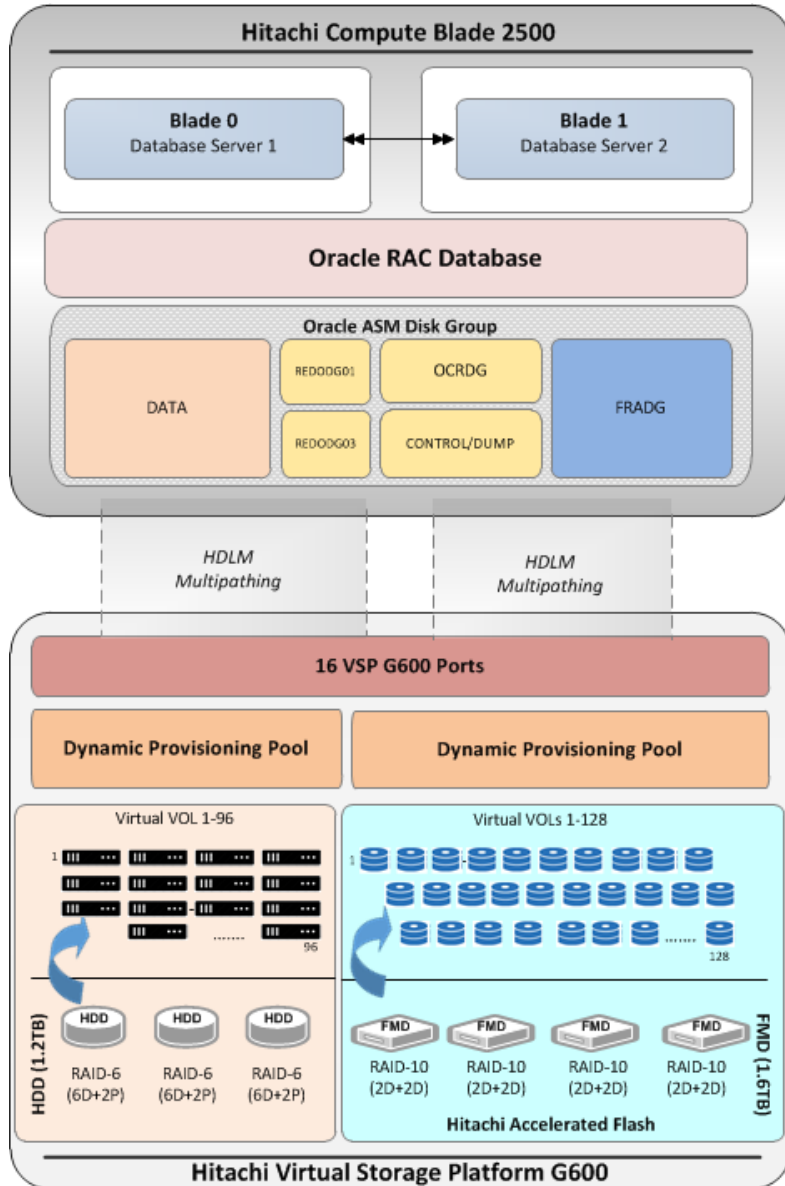


Table 3 shows a sample storage pool configuration. This environment uses Hitachi Dynamic Provisioning to create the dynamic provisioning pools.

**TABLE 3. STORAGE POOL CONFIGURATION**

<b>Pool ID</b>	Broadwell-Oracle-FMD-Pool	Broadwell-Oracle-SAS-Pool
<b>Pool Type</b>	Dynamic Provisioning	Dynamic Provisioning
<b>RAID Group</b>	1-11 – 1-14	1-1 – 1-3
<b>RAID Level</b>	RAID-10 (2D+2D)	RAID-6 (6D+2P)
<b>Number of LDEVs</b>	128	96
<b>Drive Type</b>	1.6 TB flash module drive (FMD)	1.2 TB 10k RPM SAS Drive
<b>Number of Drives</b>	16, 4 per each tray	24
<b>Pool Capacity</b>	12.80 TB	18.74 TB

Table 4 shows the logical storage configuration used in the tested configuration.

**TABLE 4. LOGICAL STORAGE CONFIGURATION**

<b>Pool ID</b>	Broadwell-Oracle-FMD-Pool	Broadwell-Oracle-SAS-Pool
<b>Number of VVOLs</b>	128	96
<b>VVOL Size</b>	102 GB	200 GB
<b>Purpose</b>	OLTP Application Tablespaces	Oracle System Sysaux Undo Temp Redo Logs Oracle Cluster Registry Voting Disk
<b>Storage Port</b>	1A, 1B, 3A, 3B, 2A, 2B, 4A, 4B, 1C, 1D, 2C, 2D, 3C, 3D, 4C, 4D	

## Database Layout

The database layout design uses recommended practices from Hitachi Data Systems for Hitachi Virtual Storage Platform G600 using Hitachi Accelerated Flash 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. 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.

- **Data and Indexes Tablespace** — Assign a Data ASM diskgroup for the data and index tablespaces. The small file table space consists of 4,096 data files that are each 8 GB. Set the tablespace to a small initial size with auto extend enabled, to maximize storage utilization.
- **TEMP Tablespace** — Place TEMP tablespace in this configuration in the CONTROL ASM diskgroup. Quite a number of small file tempfiles are created within one single small TEMP tablespace. Limit the size of each small file tempfile to 31.25 GB.
- **Undo Tablespace** — Place UNDO tablespace in this configuration in the CONTROL ASM diskgroup. Assign one BIGFILE UNDO tablespace for each database instance in a two node Oracle RAC database. Quite a number of small file undo datafiles are created within each small UNDO tablespace. Limit the size of each small undo datafile to 8 GB.
- **Online Redo Logs** — Assign one ASM diskgroup for each database instance. Four redo logs are created for each database instance in a two node Oracle RAC database. Set the size of each redo log file to 8 GB.
- **Oracle Cluster Registry and Voting Disk** — Place each of these files in this configuration in the OCRDG ASM diskgroup.
- **Size Settings** — Set the database block size to 8 KB. Set the ASM allocation unit to 4 MB.
- **ASM FILE SYSTEM I/O Settings** — Set the Oracle ASM I/O operations for database files as in Table 5.

TABLE 5. ASM DISK GROUP AND DATABASE FILE STRUCTURE

ASM Disk Group	OCRDG	CONTROL	REDO01	REDO02	FRADG	DATADG
Oracle Cluster Registry and Voting Disks	Yes					
Parameter file		Yes				
Password file		Yes				
Multiplexed Control files		Yes			Yes	
Multiplexed Redo log file			Yes	Yes		
Archive log files					Yes	
System data files		Yes				
Application small and big data files						Yes

Table 6 shows the Oracle RAC database configuration.

**TABLE 6. ORACLE RAC DATABASE SETTINGS**

For This	Use This
RAC configuration	Yes
ASM	Yes - Oracle RAC Database

Table 7 shows the Oracle environment parameters.

**TABLE 7. ORACLE ENVIRONMENT PARAMETERS**

For This	Use This
SGA_TARGET	128 GB
PGA_AGGREGATE_TARGET	64 GB
DB_CACHE_SIZE	64 GB
DB_KEEP_CACHE_SIZE	32 GB
DB_RECYCLE_CACHE_SIZE	8 GB
LOG_BUFFER	536870912
USE_LARGE_PAGES	TRUE

Table 8, “Oracle ASM Disk Configuration,” on page 12 has the details for the disk mappings from the LUNs to the operating system devices and to the ASM disk groups for Oracle RAC Database tablespaces.

**TABLE 8. ORACLE ASM DISK CONFIGURATION**

V-VOLs	LUN	LUN Size	Pool	ASM Disk	ASMDG	Purpose
00:01:2A	0	200 GB	ora_dp_sas_01	NA	N/A	Server Boot LUN (one for each server)
00:01:2B	0	200GB		NA	N/A	Server Boot LUN (one for each server)
00:01:BB - 00:01: BD	1-3	200 GB	Broadwell-Refresh- Oracle-SAS-Pool	OCRDISK1 - OCRDISK3	OCRDG	Oracle Cluster Registry and Voting Disk
00:01:6B - 00:01:7E	4-23	100 GB		REDODISK1 - REDODISK8	REDO_DG01	Online REDO Log Group
				REDODISK9 - REDODISK16	REDO_DG11	Online REDO Log Group
				SYSDISK1 - SYSDISK4	CONTROL	Sys,Undo,TEMP
00:01:7F - 00:01:C5	24-94	200 GB		FRDISK1 - FRDISK71	FRADG	Fast Recovery Area
00:02:6C - 00:02:EB	39-166	120 GB	Broadwell-Refresh- Oracle-FMD-Pool	DATADISK1 - DATADISK128	DATADG	Application Data

## Server and Application Architecture

This reference architecture uses a single Hitachi Compute Blade 2500 with two server blades. Two server blades are used for a two-node RAC configuration.


This configuration provides the compute power for 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.

**TABLE 9. SERVER DETAILS**

Server	Form Size	Server Name	Role	CPU Cores (total)	RAM	Blade
Node1	Half-Width size	RAC01	RAC Node	44	256 GB	1
Node2	Half-Width size	RAC02	RAC Node	44	256 GB	2

Figure 3 shows the server infrastructure for the reference architecture.

**Figure 3**



## SAN Architecture

Map the provisioned LDEVs to multiple ports on Hitachi Virtual Storage Platform G600 using Hitachi Accelerated Flash. These LDEV port assignments provide multiple paths to the storage system from the host for high availability.

Each of the database servers uses four Fibre Channel ports, with two ports from each of the PCIe HBA cards from Hitachi listed in Table 1, “Key Hardware Components,” on page 3. This provides a four path connection for all LUNs mapped to each of the database servers in Oracle RAC database. Table 10 shows connections from the HBAs to the Hitachi Virtual Storage Platform G600 storage ports.

**TABLE 10. FIBRE CHANNEL CONNECTIONS TO HITACHI VIRTUAL STORAGE PLATFORM G600**

Server	Host	HBA Ports	Storage Port	Storage Host Group	Zone Names
Server Blade 1	RAC-1	HBA1-1	CL1-A	CB2500_229_B0_HBA1_1	CB2500_229_B0_HBA1_1_ASE42.33_1A_3A
			CL3-A		
		HBA1-2	CL1-B	CB2500_229_B0_HBA1_2	CB2500_229_B0_HBA1_2_ASE42.33_1B_3B
			CL3-B		
		HBA2-1	CL2-A	CB2500_229_B0_HBA2_1	CB2500_229_B0_HBA2_1_ASE42.33_2A_4A
			CL4-A		
		HBA2-2	CL2-B	CB2500_229_B0_HBA2_2	CB2500_229_B0_HBA1_2_ASE42.33_2B_4B
			CL4-B		



**TABLE 10. FIBRE CHANNEL CONNECTIONS TO HITACHI VIRTUAL STORAGE PLATFORM G600 (CONTINUED)**


Server	Host	HBA Ports	Storage Port	Storage Host Group	Zone Names
Server Blade 2	RAC-2	HBA1-1	CL1-C	CB2500_229_B1_HBA1_1	CB2500_229_B1_HBA1_1_ ASE42.33_1C_3C
			CL3-C		
		HBA1-2	CL1-D	CB2500_229_B1_HBA1_2	CB2500_229_B1_HBA1_1_ ASE42.33_1D_3D
			CL3-D		
		HBA2-1	CL2-C	CB2500_229_B1_HBA2_1	CB2500_229_B1_HBA1_1_ ASE42.33_2C_4C
			CL4-C		
		HBA2-2	CL2-D	CB2500_229_B1_HBA2_2	CB2500_229_B1_HBA1_1_ ASE42.33_2D_4D
			CL4-D		

*Hitachi 16 Gb/sec PCIe HBA Card Configuration*

This describes the configuration for the Hitachi 16 Gb/sec PCIe HBA cards that are used on the server blades.

Figure 4 shows the Hitachi 16 Gb/sec HBA PCIe cards that are installed in the Hitachi Compute Blade 2500 chassis.

**Figure 4**



Set the following parameter for each of the Hitachi HBA PCIe cards following Table 11.

**TABLE 11. HITACHI HBA PCIE CARD PARAMETERS**

For This	Set This
Boot Function	Enable
Link Speed	16Gbps
Connection Type	Point-to-Pont
Multiple Port ID	Disable
Select Boot Device	Enable
Multipath Function	Enable

## 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 cache synchronization of Oracle RAC and inter-node communication amongst the nodes in the cluster.
- **Public Network** — This network provides client connections to the applications and Oracle RAC.

Hitachi Data Systems recommends using a pair of 10 Gb/sec NICs for the cluster interconnect and public network.

Each server blade in this reference architecture has a quad port 10 Gb/sec onboard NIC. The NIC ports have interconnected links to the two internal 10 Gb/sec Ethernet switches in the chassis.

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

- For each server in the Oracle RAC clusterware configuration, use at least two identical, high bandwidth, low-latency NICs for the interconnection.
- Use NIC bonding to provide fail over 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.

Figure 5 shows the network configuration for the reference architecture environment.

**Figure 5**




Table 12 shows the network configuration for this solution. Configure the VLAN to fit your network environment.

**TABLE 12. NETWORK CONFIGURATION**

Server	NIC Ports	UMC PF Number	NIC Bond	VLAN	IP Address	Network
Server Blade 1	B0-NIC-0	0	Bond1	168	192.168.109.10	Private
	B0-NIC-1	1				Private
	B0-NIC-2	2	Bond2	242	172.17.242.70	Public
		6	Bond3	244	172.17.244.70	Management
	B0-NIC-3	3	Bond2	242	172.17.242.70	Public
		7	Bond3	244	172.17.244.70	Management
Server Blade 2	B1-NIC-0	0	Bond1	168	192.168.109.11	Private
	B1-NIC-1	1				Private
	B1-NIC-2	2	Bond2	242	172.17.242.71	Public
		6	Bond3	244	172.17.244.71	Management
	B1-NIC-3	3	Bond2	242	172.17.242.71	Public
		7	Bond3	244	172.17.244.71	Management

## Engineering Validation

This summarizes the key observations from the test results for Hitachi Unified Compute Platform 6000 for Oracle Real Application Cluster using Hitachi Virtual Storage Platform G600 with Hitachi Accelerated Flash. The test results come from using these 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](#).

- **peakmarks**

[peakmarks](#) is the leading benchmark software for Oracle platforms 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 and cloud services.

## Test Methodology

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

**TABLE 13. ORACLE ORION AND PEAKMARKS TEST RESULTS**

Test Case	Tool	Test/Workload Type	Measurement/Unit	Results with 16 FMD
1	ORION	OLTP Random Read	Max. IO/s	<b>380,057</b>
			Avg. RT	<b>0.97</b>
2	ORION	OLTP Random Writes	Max. IO/s	<b>118,737</b>
			Avg. RT	<b>1.2</b>
3	ORION	OLAP Sequential Reads	Max.Throughput (GB/s)	<b>11.13</b>
4	ORION	OLAP Sequential Writes	Max.Throughput (GB/s)	<b>4.71</b>
7	peakmarks	1 node – OLTP Random Read	Max. IO/s	<b>195,900</b>
			Avg. RT	<b>0.8</b>
8	peakmarks	1 node – OLTP Random Writes	Max. IO/s	<b>172,100</b>
			Avg. RT	<b>4.9</b>
9	peakmarks	1 node – OLAP Sequential Reads	Max.Throughput (GB/s)	<b>6.2</b>

**TABLE 13. ORACLE ORION AND PEAKMARKS TEST RESULTS (CONTINUED)**

Test Case	Tool	Test/Workload Type	Measurement/Unit	Results with 16 FMD
10	peakmarks	2 nodes – OLTP Random Read	Max. IO/s	<b>266,400</b>
			Avg. RT	<b>1.4</b>
11	peakmarks	2 nodes – OLTP Random Writes	Max. IO/s	<b>173,300</b>
			Avg. RT	<b>2.4</b>
12	peakmarks	2 nodes – OLAP Sequential Reads	Max.Throughput (GB/s)	<b>6.3</b>

## Test Results

Each table defined with different ORION I/O performance results and comparison.

**TABLE 14. VSP G600, 1.6 TB FMD, RAID LEVEL SCALABILITY, RAID-10 (2D+2D), 100% SEQUENTIAL WRITES, 1024 KB BLOCK SIZE, 8 THREADS PER FMD, MAX I/O RATE**

Total FMDs	Total LDEVs	Total V-VOLs	GB/s
16	128	106	4.71

**TABLE 15. VSP G600, 1.6 TB FMD, RAID LEVEL SCALABILITY, RAID-10 (2D+2D), 100% SEQUENTIAL READ, 1024 KB BLOCK SIZE, 8 THREADS PER FMD, MAX I/O RATE**

Total FMDs	Total LDEVs	Total V-VOLs	GB/s
16	128	106	11.13

**TABLE 16. VSP G600, 1.6 TB FMD, RAID LEVEL SCALABILITY, RAID-10 (2D+2D), 100% RANDOM READ, 8 KB BLOCK SIZE, 8 THREADS PER FMD, MAX I/O RATE**

Total FMDs	Total LDEVs	Total V-VOLs	IOPS	RT (ms)
16	128	106	380,057	0.97

**TABLE 17. VSP G600, 1.6 TB FMD, RAID LEVEL SCALABILITY, RAID-10 (2D+2D), 100% RANDOM WRITES, 8 KB BLOCK SIZE, 8 THREADS PER FMD, MAX I/O RATE**

Total FMDs	Total LDEVs	Total V-VOLs	IOPS	RT (ms)
16	128	106	118,737	1.2

**TABLE 18. VSP G600, 1.6 TB FMD, RAID LEVEL SCALABILITY, RAID-10 (2D+2D), 75% RANDOM READS, 256 KB BLOCK SIZE, 8 THREADS PER FMD, MAX I/O RATE**

Total FMDs	Total LDEVs	Total V-VOLs	IOPS
16	128	106	230,103

## For More Information

Hitachi Data Systems 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.

Live and recorded product demonstrations are available for many Hitachi products. To schedule a live demonstration, contact a sales representative. To view a recorded demonstration, see the [Resources](#) website.

Hitachi Data Systems Academy provides best-in-class training on Hitachi products, technology, solutions and certifications. Hitachi Data Systems Academy delivers on-demand web-based training (WBT), classroom-based instructor-led training (ILT) and virtual instructor-led training (vILT) courses. For more information, see the Hitachi Data Systems Services [Training and Certification](#) website.

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

Corporate Headquarters  
2845 Lafayette Street  
Santa Clara, CA 95050-2639 USA  
[www.HDS.com](http://www.HDS.com)    [community.HDS.com](http://community.HDS.com)

Regional Contact Information  
**Americas:** +1 866 374 5822 or [info@hds.com](mailto:info@hds.com)  
**Europe, Middle East and Africa:** +44 (0) 1753 618000 or [info.emea@hds.com](mailto:info.emea@hds.com)  
**Asia Pacific:** +852 3189 7900 or [hds.marketing.apac@hds.com](mailto:hds.marketing.apac@hds.com)

© Hitachi Data Systems Corporation, 2017. All rights reserved. HITACHI is a trademark or registered trademark of Hitachi, Ltd. VSP is a trademark or registered trademark of Hitachi Data Systems 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 Data Systems Corporation.