HITACHI
Inspire the Next

Hitachi Unified Compute Platform 6000 for Oracle Real Application Clusters on Two Nodes Using Hitachi Virtual Storage Platform G600, Hitachi Accelerated Flash, and Hitachi Compute Blade 2500 With Intel Xeon E7-8893 v3 Processor CPUs

TECHNICAL PAPER

Reference Architecture Guide

By Ben Prusinski; Steve Burns

May 2016

Hitachi Data Systems

Feedback

Hitachi Data Systems welcomes your feedback. Please share your thoughts by sending an email message to 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.

Contents

Hitachi Unified Compute Platform 6000 for Oracle Real Application Clusters on Two Nodes Using Hitachi Virtual Storage Platform G600, Hitachi Accelerated Flash, and Hitachi Compute Blade 2500 With Intel Xeon E7-8893 v3 Processor CPUs

Reference Architecture Guide

This reference architecture guide shows how using Hitachi Unified Compute Platform 6000 for Oracle Real Application Clusters (UCP for Oracle RAC) provides a high performance, integrated solution for an Oracle infrastructure. The environment uses Hitachi Virtual Storage Platform G600 (VSP G600) with Hitachi Accelerated Flash (HAF). Use this document to design an infrastructure for your requirements and budget.

This validated solution integrates servers, storage systems, network, and storage software. The environment provides reliability, high availability, scalability, and performance while processing small-scale to large-scale OLTP workloads. The dedicated server runs Oracle Database 12c Release 1 with the Oracle Real Application Clusters (RAC) option. The operating system is Red Hat Enterprise Linux 6.6.

This reference architecture document is for the following roles:

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

To use this reference architecture guide, familiarity with the following is required:

- Hitachi Virtual Storage Platform G600 using Hitachi Accelerated Flash
- Hitachi Compute Blade 2500 (CB 2500) using 520X B2 server blades with Intel Xeon E7-8893 v3 processor CPUs
- Storage area networks
- Oracle RAC Database 12c Release 1
- Oracle Automatic Storage Management (Oracle ASM)
- Red Hat Enterprise Linux (RHEL)

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 Clusters on two nodes using Hitachi Virtual Storage Platform G600 with Hitachi Accelerated Flash. In this configuration, Hitachi Compute Blade 2500 uses 520X B2 server blades with Intel Xeon E7-8893 v3 processor CPUs.

This environment addresses the high availability, performance, and scalability requirements for on-line transaction processing (OLTP) workloads. Tailor your implementation of this solution to meet your specific needs.

This reference architecture includes the following:

- **Hitachi Compute Blade 2500** chassis with three server blades
  - **Server Blade 1** — Oracle RAC NODE 1
  - **Server Blade 3** — Oracle RAC NODE 2
  - **Server Blade 5** — N+M cold standby node
- **Hitachi Virtual Storage Platform G600** with Hitachi Accelerated Flash
- 16 Gb/sec direct-connect SAN infrastructure
- 10 GbE LAN infrastructure

Figure 1 on page 3 shows the high level infrastructure for this solution.

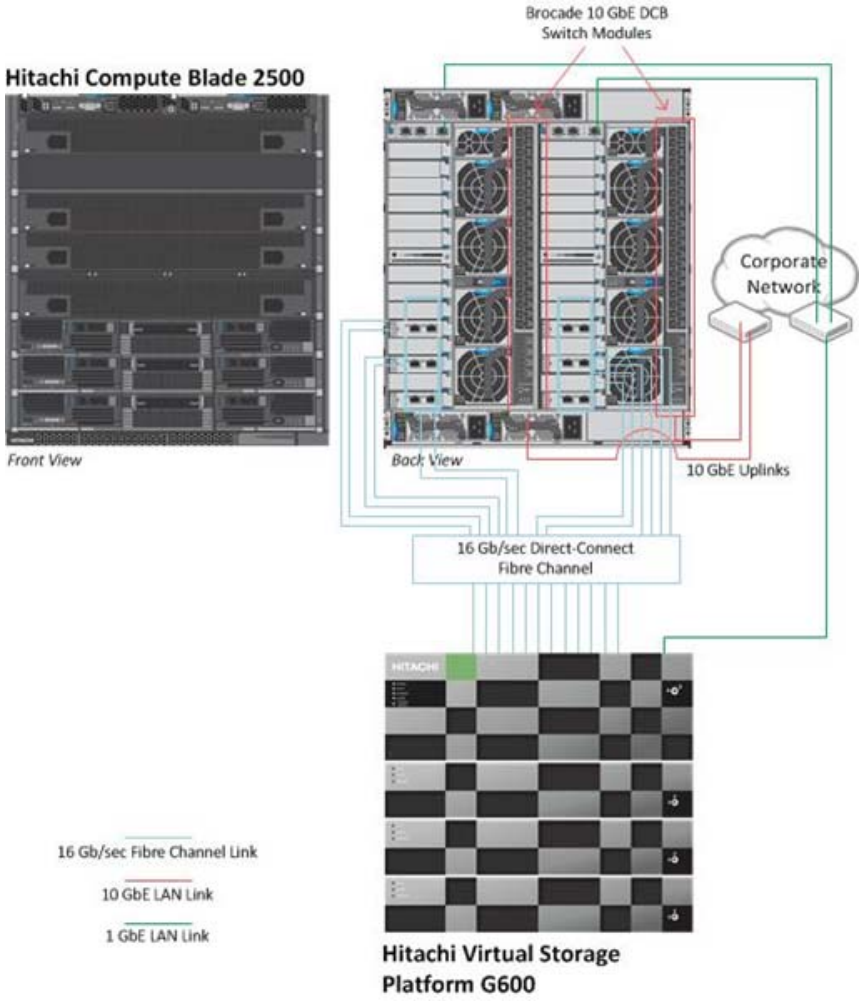
Brocade 10 GbE DCB
Switch Modules
Hitachi Compute Blade 2500
Front View
Back View
Corporate
Network
10 GbE Uplinks
16 Gb/sec Direct-Connect
Fibre Channel
16 Gb/sec Fibre Channel Link
10 GbE LAN Link
1 GbE LAN Link
Hitachi Virtual Storage
Platform G600

Figure 1

Key Solution Components

Table 1 lists the key hardware components used in this reference architecture.

Table 1. Key Solution Components From Hitachi Data System

Hardware	Detail Description	Version	Quantity
Hitachi Virtual Storage Platform G600	■ Dual controller ■ 16 × 16 Gb/sec Fibre Channel ports ■ 16 back-end serial attached SCSI (SAS) ports ■ 256 GB cache memory	83-01-03-40/00	1
	■ 1.6 TB flash memory drives (FMDs)		16
	■ 1.2 TB 10k RPM SAS drives		24
Hitachi Compute Blade 2500 chassis	■ 2 × 10 GbE data center bridging (DCB) local area network (LAN) switch module ■ 10 Fan modules ■ 2 Management modules ■ 4 Power supply modules	Management Module Firmware Version A0122-B-1009 Dictionary Version A0013 DCB Switch Version 4.0.1_hit1	1
520X B2 full width server blade	■ 2 Intel Xeon E7-8893 v3 processor CPUs ■ 256 GB RAM, 16 × 16 GB DDR4 memory ■ 1 × 4-port 10 GbE converged network adapter (CNA) LAN on motherboard (LOM)	09-07	3
	■ Hitachi 16 Gb/sec 2-port PCIe Fibre Channel HBA	40-03-07	6

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

Table 2. Key Software Components

Software	Version	Function
Hitachi Storage Virtualization Operating System with the following: ■ Hitachi Storage Navigator ■ Hitachi Dynamic Tiering ■ Hitachi Dynamic Provisioning	N/A	Storage management suite
Hitachi Command Suite	8	Storage management suite
Hitachi Compute Systems Manager	8.1.4-04	N+M Management
Red Hat Enterprise Linux	6.6	Operating system for Oracle RAC

Table 2. Key Software Components (Continued)

Software	Version	Function
Oracle Grid Infrastructure	12c Release 1	Oracle Clusterware
Oracle Database	12c Release 1	Oracle database system
Hitachi Dynamic Link Manager Advanced	8.1.4-00	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 uses 520X B2 full width server blades. This provides the scalability and flexibility for an Oracle RAC configuration with cold standby server using built using N+M server 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. New management software improves ease of use to save time and reduce complexity. The infrastructure of Storage Virtualization Operating System creates a management framework for improved IT response to business demands.

This solution uses Hitachi Virtual Storage Platform G600, which 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.

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 spans and integrates multiple platforms. It is 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 base. This provides storage configuration and control capabilities for you.

Storage Virtualization Operating System uses Hitachi Dynamic Provisioning 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. This avoids intermixing different type of data I/O. which benefits performance during peak database operations.

This solution uses Hitachi Dynamic Tiering, a part of Storage Virtualization Operating System. Separately licensed, Dynamic Tiering virtualizes and automates mobility between tiers for maximum performance and efficiency.

The dynamic tiering pool used for this solution provides you the capability to introduce different types of disks at a later time when necessary to increase the disk pool capacity. Maintain performance when accessing frequently used data for Oracle on flash modules while placing less frequently accessed data on other types of disks.

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 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.

Based on test results, the extended round robin policy algorithm for Dynamic Link Manager Advanced provides the best overall performance with Oracle environments.

Brocade Switches

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 uses the following Brocade products:

- Brocade 10 GbE switch module

Red Hat Enterprise Linux

Using the stability and flexibility of Red Hat Enterprise Linux, reallocate your resources towards meeting the next challenges instead of maintaining the status quo. Deliver meaningful business results by providing exceptional reliability of military-grade security. Use Enterprise Linux to tailor your infrastructure as markets shift and technologies evolve.

Oracle Database With Oracle Real Application Clusters

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.

Oracle Clusterware is portable cluster software that allows clustering of independent servers so that they cooperate as a single system. Oracle Clusterware is the required cluster technology for Oracle Real Application Clusters.

Solution Design

This section describes the reference architecture environment, implementing a quarter-rack environment for Hitachi Unified Compute Platform 6000 for Oracle Real Application Clusters using Hitachi Virtual Storage Platform G600 with Hitachi Accelerated Flash.

Specific infrastructure configuration includes the following:

- **Server** — Configure two server nodes in an Oracle Real Application Cluster with one additional N+M standby node.
- **Storage System** — Map the LDEVs to each port that is presented to the server as LUNs.
- **SAN Connection** — Connect each Fibre Channel HBA port directly to the storage front-end ports.

Storage Architecture

This section 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 Configuration

This is the high level storage configuration of this solution.

Figure 2 on page 9 describes the storage configuration for this solution.

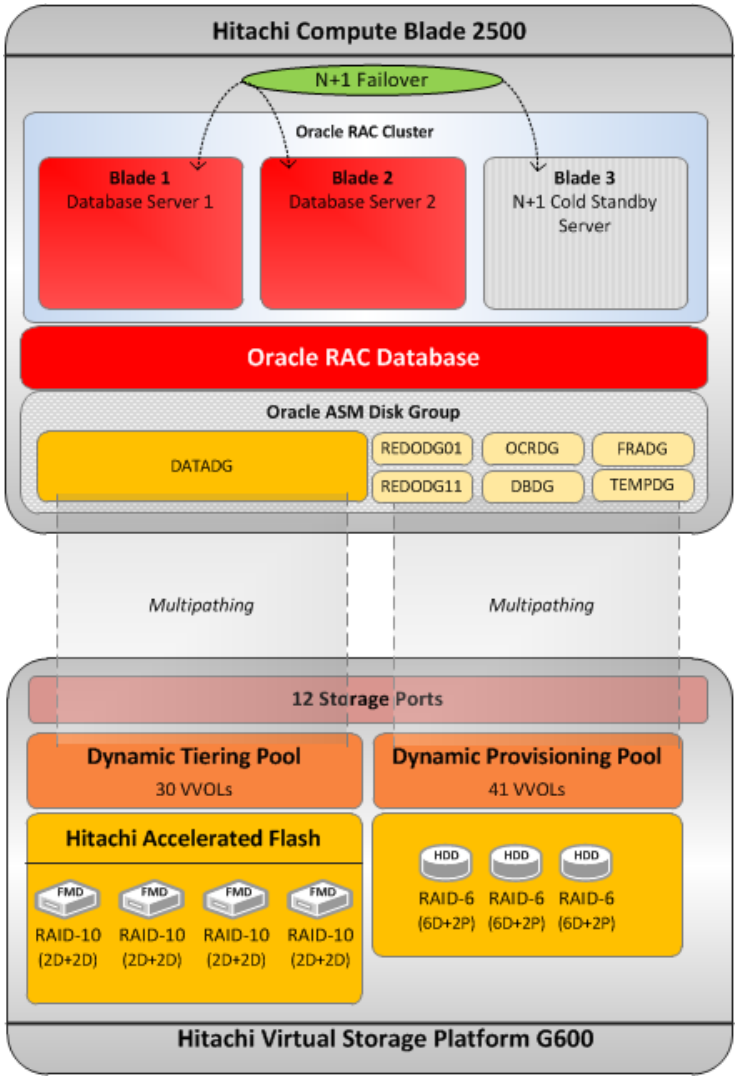
Hitachi Compute Blade 2500
N+1 Failover
Oracle RAC Cluster
Blade 1
Database Server 1
Blade 2
Database Server 2
Blade 3
N+1 Cold Standby
Server
Oracle RAC Database
Oracle ASM Disk Group
DATADG
REDODG01
REDODG11
OCRDG
DBDG
FRADG
TEMPDG
Multipathing
Multipathing
12 Storage Ports
Dynamic Tiering Pool
30 VVOLs
Dynamic Provisioning Pool
41 VVOLs
Hitachi Accelerated Flash
FMD
RAID-10
(2D+2D)
HDD
RAID-6
(6D+2P)
Hitachi Virtual Storage Platform G600

Figure 2

Table 3 shows the storage pool configuration used in the tested configuration.

Table 3. Storage Pool Configuration

Pool ID	ora_dp_fmd_01	ora_dp_sas__01
Pool Type	Dynamic Tiering	Dynamic Provisioning
RAID Group	1-1 – 1-4	1-5 – 1-7
RAID Level	RAID-10 (2D+2D)	RAID-6 (6D+2P)
Number of LDEVs	4	6
Drive Type	1.6 TB Flash Module Drive (FMD)	1.2 TB 10k RPM SAS Drive
Number of Drives	16, 4 per each tray	24
Pool Capacity	11.99 TB	17.93 TB

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

Table 4. Logical Storage Configuration

Pool ID	ora_dp_fmd_01	ora_dp_sas_01
Number of VVOLs	30	41
VVOL Size	400 GB	200 GB
Purpose	OLTP Application Tablespaces	■ System ■ Sysaux ■ Undo ■ Temp ■ Redo Logs ■ Oracle Cluster Registry ■ Voting Disk
Storage Port	1A, 1B, 3A, 3B, 2A, 2B, 4A, 4B	

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 an ASM diskgroup for the data and index tablespaces. The small file table space consists of 364 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 Data ASM diskgroup. Quite a number of small file tempfiles are created within one single small TEMP tablespace.
- **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** — 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** — 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.
- **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 follows:

  FILESYSTEMIO_OPTIONS = setall

- **ASM Disk I/O parameter** — Set the ASM parameter to TRUE as shown below:

  Disk_asynch_io-boolean = TRUE

Table 5 shows the Oracle RAC database configuration.

Table 5. Oracle RAC Database Settings

Environment	Value
RAC configuration	Yes
ASM	Yes - Oracle RAC Database

Table 6 shows the Oracle Environment Parameters.

Table 6. Oracle Environment Parameters

Setting	Parameter
SGA_TARGET	70 GB
PGA_AGGREGATE_TARGET	14 GB
DB_CACHE_SIZE	16 GB
DB_KEEP_CACHE_SIZE	16 GB
DB_RECYCLE_CACHE_SIZE	16 GB
LOG_BUFFER	256 MB
USE_LARGE_PAGES	TRUE

Table 7 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 7. Oracle ASM Disk Configuration

ASM Disk Group	ASM Disk	HDLM LUNs	LUNs Count	Purpose
OCRDG	OCRDISK1	/dev/sddlmbo	3	Oracle Cluster Registry and Voting Disk
	OCRDISK2	/dev/sddlmbp		
	OCRDISK3	/dev/sddlmca		
REDO_DG01	REDODISK1	/dev/sddlmcb	4	Online REDO log group and control file
	REDODISK2	/dev/sddlmcc		
	REDODISK3	/dev/sddlmcd		
	REDODISK4	/dev/sddlmce		
REDO_DG11	REDODISK5	/dev/sddlmcf	4	
	REDODISK6	/dev/sddlmg		
	REDODISK7	/dev/sddlmch		
	REDODISK8	/dev/sddlmci		

Table 7. Oracle ASM Disk Configuration (Continued)

ASM Disk Group	ASM Disk	HDLM LUNs	LUNs Count	Purpose
FRADG	FRADISK1	/dev/sddlmcj	16	Flash Recovery Area
	FRADISK2	/dev/sddlmck		
	FRADISK3	/dev/sddlmcl		
	FRADISK4	/dev/sddlmcm		
	FRADISK5	/dev/sddlmcn		
	FRADISK6	/dev/sddlmco		
	FRADISK7	/dev/sddlmcp		
	FRADISK8	/dev/sddlmda		
	FRADISK9	/dev/sddlmdb		
	FRADISK10	/dev/sddlmdc		
	FRADISK11	/dev/sddlmdd		
	FRADISK12	/dev/sddlmde		
	FRADISK13	/dev/sddlmdf		
	FRADISK14	/dev/sddlmdg		
	FRADISK15	/dev/sddlmdh		
	FRADISK16	/dev/sddlmdi		
TEMPDG	TEMPDISK1	/dev/sddlmdj	8	Temp
	TEMPDISK2	/dev/sddlmdk		
	TEMPDISK3	/dev/sddlmdl		
	TEMPDISK4	/dev/sddlmdm		
	TEMPDISK5	/dev/sddlmdn		
	TEMPDISK6	/dev/sddlmdo		
	TEMPDISK7	/dev/sddlmdp		
	TEMPDISK8	/dev/sddlmea		
DBDG	SYSDISK1	/dev/sddlmeb	6	Sys, Undo
	SYSDISK2	/dev/sddlmec		
	SYSDISK3	/dev/sddlmed		
	SYSDISK4	/dev/sddlmee		
	SYSDISK5	/dev/sddlmef		
	SYSDISK6	/dev/sddlmeg		

Table 7. Oracle ASM Disk Configuration (Continued)

ASM Disk Group	ASM Disk	HDLM LUNs	LUNs Count	Purpose
DATADG	DATADISK1	/dev/sddlmaa	30	Application Data
	DATADISK2	/dev/sddlmab		
	DATADISK3	/dev/sddlmac		
	DATADISK4	/dev/sddlmad		
	DATADISK5	/dev/sddlmae		
	DATADISK6	/dev/sddlmaf		
	DATADISK7	/dev/sddlmag		
	DATADISK8	/dev/sddlmah		
	DATADISK9	/dev/sddlmai		
	DATADISK10	/dev/sddlmaj		
	DATADISK11	/dev/sddlmak		
	DATADISK12	/dev/sddlmal		
	DATADISK13	/dev/sddlmam		
	DATADISK14	/dev/sddlman		
	DATADISK15	/dev/sddlmao		
	DATADISK16	/dev/sddlmap		
	DATADISK17	/dev/sddlmba		
	DATADISK18	/dev/sddlmbb		
	DATADISK19	/dev/sddlmbc		
	DATADISK20	/dev/sddlmbd		
	DATADISK21	/dev/sddlmbe		
	DATADISK22	/dev/sddlmbf		
	DATADISK23	/dev/sddlmbg		
	DATADISK24	/dev/sddlmbh		
	DATADISK25	/dev/sddlmbi		
	DATADISK26	/dev/sddlmbj		
	DATADISK27	/dev/sddlmbk		
	DATADISK28	/dev/sddlmbl		
	DATADISK29	/dev/sddlmbm		
	DATADISK30	/dev/sddlmbn		

Server and Application Architecture

This reference architecture uses a single Hitachi Compute Blade 2500 chassis with three server blades. Two server blades are used for a two-node Oracle RAC configuration with an additional N+M cold standby node.

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.

Table 8. Server Details

Server	Form Size	Server Name	Role	CPU Core	RAM	Server Blade
Node1	Full Width	RAC01	RAC Node	8	256 GB	1
Node2	Full Width	RAC02	RAC Node	8	256 GB	3
Node3	Full Width	N+M	Standby	8	256 GB	5

N+M Cold Standby Server

This solution uses an N+M cold standby server on Hitachi Compute Blade 2500 to provide server redundancy. N+M cold standby provides automated server blade fault detection and failover in the event of a server blade failure. Server Blade 3 is configured as the cold standby node for the two active nodes on Server Blade 1 and Server Blade 2.

Note — Make sure to enable N+M cold standby before you configure the storage host group that is part of the N+M setup. The N+M configuration enables an additional virtual WWN that is required for the N+M setup.

Figure 3 shows the server infrastructure for the reference architecture with the cold standby node.

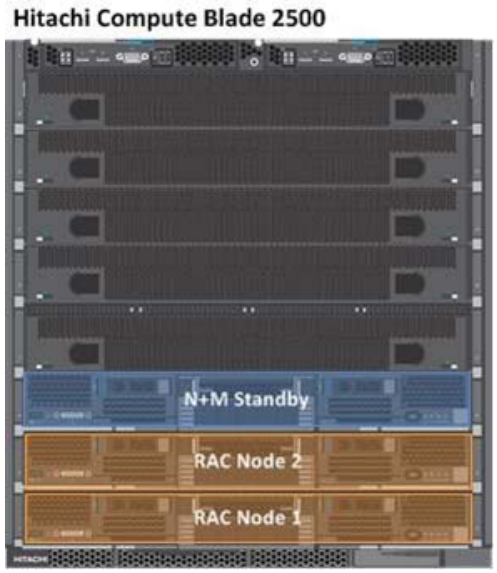

Figure 3

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 Solution Components From Hitachi Data System,” on page 4. This provides a four path connection for all LUNs mapped to each of the database servers in the Oracle RAC database. Table 9 shows the direct-connect from the HBA of the server blade to the Hitachi Virtual Storage Platform G600 ports.

Table 9. Fibre Channel Direct-Connect Configuration on Hitachi Virtual Storage Platform G600

Chassis	Host	HBA	Storage Port	Storage Host Group
Hitachi Compute Blade 2500	BLADE 1	HBA1-1	1A	BS25K_B1_HBA1_1
		HBA1-2	2A	BS25K_B1_HBA2_2
		HBA2-1	1B	BS25K_B1_HBA2_1
		HBA2-2	2B	BS25K_B1_HBA2_2
	BLADE 3	HBA1-1	3A	BS25K_B3_HBA1_1
		HBA1-2	4A	BS25K_B3_HBA1_2
		HBA2-1	3B	BS25K_B3_HBA2_1
		HBA2-2	4B	BS25K_B3_HBA2_2
	BLADE 5	HBA1-1	1C	BS25K_B1_HBA1_1 BS25K_B3_HBA1_1
		HBA1-2	2C	BS25K_B1_HBA1_2 BS25K_B3_HBA1_2
		HBA2-1	1D	BS25K_B1_HBA2_1 BS25K_B3_HBA2_1
		HBA2-2	2D	BS25K_B1_HBA2_2 BS25K_B3_HBA2_2

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.

Note — Make sure that the N+M feature is enabled before you configure the storage host group that is part of the N+M setup. The N+M configuration enables additional virtual WWN that is required for the N+M setup.

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

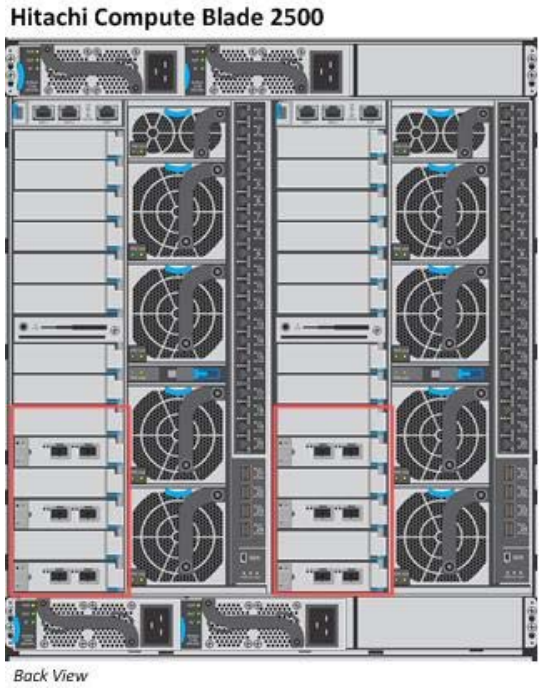

Figure 4

Set the following parameter for each of the Hitachi HBA PCIe cards following Table 10.

Table 10. Hitachi HBA PCIe Card Parameters

Setting	Parameter
Boot Function	**Enable**
Link Speed	**16Gbps**
Connection Type	**Point-to-Point**
Multiple Port ID	**Enable**
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 GbE NICs for the cluster interconnect and public network.

Each server blade in this reference architecture has a quad port 10 GbE onboard NIC. The NIC ports have interconnected links to the two internal 10 GbE 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.

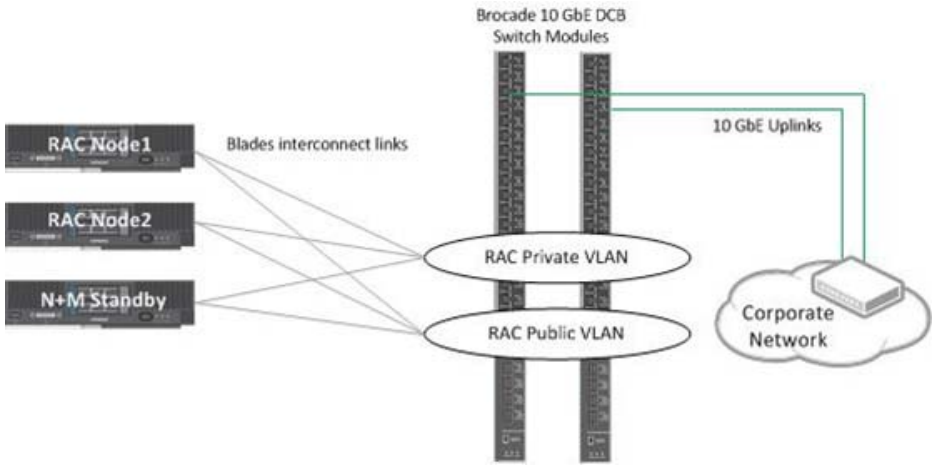

Figure 5

Table 11 shows the network configuration for this solution. Configure the VLAN to fit your network environment.

Table 11. Network Configuration

Server	NIC Ports	UMC PF#	Switch Bay ID	Switch Ports (Internal)	VLAN/Subnet	NIC BOND	Network	Bandwidth (Gb/sec)
Database Server 1	B1-CNIC-0	0	1	1	208	Bond1	Private	10
	B1-CNIC-1	1	2	1		Available		10
	B1-CNIC-2	2	1	2	243	Bond200	Public Oracle	9
		6	1	2	244	Bond300	Public Management	1
	B1-CNIC-3	3	2	2	243	Bond200	Public Oracle	9
		7	2	2	244	Bond300	Public Management	1
Database Server 2	B2-CNIC-0	0	1	3	208	Bond1	Private	10
	B2-CNIC-1	1	2	3		Available		10
	B2-CNIC-2	2	1	4	243	Bond200	Public	9
		6	1	4		Bond300		1
	B2-CNIC-3	3	2	4		Bond200		9
		7	2	4		Bond300		1
N+1 Cold Standby Server	B2-CNIC-0	0	1	5	208	From failed server	Private	10
	B2-CNIC-1	1	2	5				10
	B2-CNIC-2	2	1	6			Public	9
		6	1	6				1
	B2-CNIC-3	3	2	6				9
		7	2	6				1

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 Hitachi Data Systems Global 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 Hitachi Data Systems Corporate Resources website. Click the **Product Demos** tab for a list of available recorded demonstrations.

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 Education website.

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

Hitachi Data Systems

Corporate Headquarters
2845 Lafayette Street
Santa Clara, CA 96050-2639 USA
www.HDS.com community.HDS.com

Regional Contact Information
Americas: +1 408 970 1000 or info@hds.com
Europe, Middle East and Africa: +44 (0) 1753 618000 or info.emea@hds.com
Asia Pacific: +852 3189 7900 or hds.marketing.apac@hds.com

AS-484-02. May 2016.