ORACLE Diamond Partner

HITACHI Inspire the Next

Hitachi Unified Compute Platform 6000 for Oracle Real Application Cluster on Two Nodes Using Hitachi Virtual Storage Platform G600 with Hitachi Accelerated Flash

TECHNICAL PAPER

Reference Architecture Guide

By Nathan Tran, Steve Burns, Kishore Daggubati

September 2015

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 Cluster on Two Nodes Using Hitachi Virtual Storage Platform G600 with Hitachi Accelerated Flash

Reference Architecture Guide

This reference architecture guide shows how using Hitachi Unified Compute Platform 6000 for Oracle Real Application Cluster option provides a high performance, integrated solution for an Oracle infrastructure. The environment uses Hitachi Virtual Storage Platform G600 with Hitachi Accelerated Flash. 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 R1 with the Oracle Real Application Cluster (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
- Hitachi network attached storage
- Storage area networks
- Oracle RAC Database 12c Release 1
- Oracle Automatic Storage Management
- Red Hat Enterprise Linux

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 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 four server blades
  - **Server Blade 1** — Oracle RAC NODE 1
  - **Server Blade 2** — Oracle RAC NODE 2
  - **Server Blade 3** — N+M cold standby node
  - **Server Blade 14** — Oracle RMAN backup server
- **Hitachi Virtual Storage Platform G600** with Hitachi Accelerated Flash
- **Hitachi NAS Platform 4060**
- 16 Gb/sec direct-connect SAN infrastructure
- 10 GbE LAN infrastructure

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

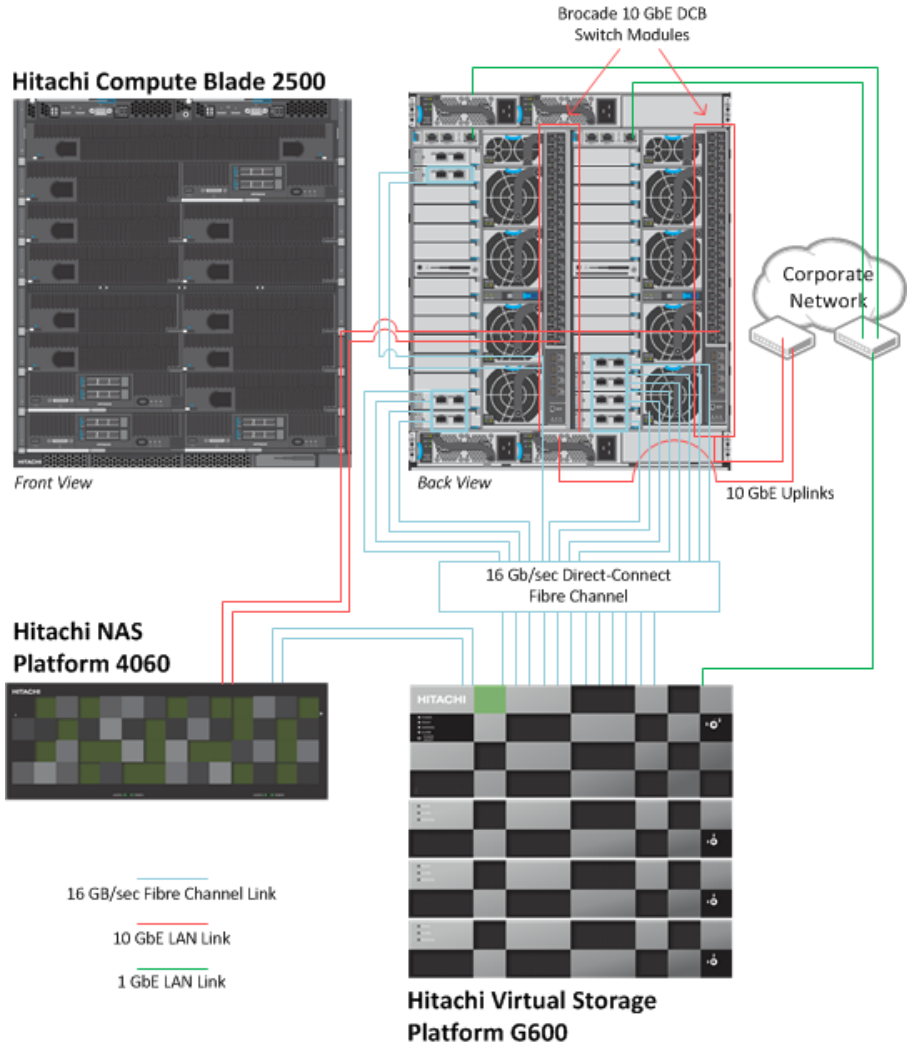

Figure 1

The tested configuration uses the same Hitachi Virtual Storage Platform G600 that hosted the production database to store of backup data from Hitachi HNAS Platform 4060. To avoid any performance impact to the production database, Hitachi Data Systems recommends using a configuration with a dedicated storage system for the production database and another storage system for storing backup data.

Key Solution Components

Table 1 list 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 backend serial attached SCSI (SAS) ports 256 GB cache memory	83-01-03-40/00	1
	1.6 TB flash memory drives (FMDs) ▪ 1 spare		17
	1.2 TB 10k RPM SAS Drives ▪ 3 spares		59
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
520H B3 Half Server Blade	2 Intel Xeon E5-2699 v3 CPU 256 RAM, 16 x 16 GB DDR4 memory 1 × 4-port 10 GbE converged network adapter (CNA) LAN on motherboard (LOM)	08-30	4
	Hitachi 16 Gb/sec 2-port PCIe Fibre Channel HBA	40-03-07	8
Hitachi HNAS Platform 4060	2 × 10 Gb/sec Cluster interconnect ports 2 × 10 Gb/sec Ethernet ports 6 × 1 Gb/sec Ethernet ports 4 × 8 Gb/sec Fibre Channel ports 2 Ethernet management ports 100 TB of capacity for data backup	12.4	1

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

Table 2. Key Software Components

Software	Version	Function
Hitachi Storage Navigator	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
Oracle ASM	12c Release 1	Oracle ASM
Oracle Database	12c Release 1	Oracle database system
Oracle RMAN	12c Release 1	Database backup
Hitachi Dynamic Link Manager Advanced	8.1.4-00	Multipath software
Hitachi Dynamic Tiering	Virtual Storage Platform G600	Storage license
Hitachi Dynamic Provisioning	Virtual Storage Platform G600	Storage license
Flash optimization feature	Virtual Storage Platform G600	Storage license
Hitachi Accelerated Flash	Virtual Storage Platform G600	Storage license
Hitachi NAS Platform firmware	12.4	NAS Platform

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.

Hitachi Compute Blade 2500 with 520HB3 Half Blade provides the scalability and flexibility for an Oracle RAC configuration with cold standby server using built in N+M server redundancy.

Hitachi Virtual Storage Platform G600

Hitachi Virtual Storage Platform G600 provides an always-available, agile, and automated foundation that you need for a continuous infrastructure cloud. This delivers enterprise-ready software-defined storage, advanced global storage virtualization, and powerful storage.

Supporting always-on operations, Virtual Storage Platform G600 includes self-service, non-disruptive migration and active-active storage clustering for zero recovery time objectives. Automate your operations with self-optimizing, policy-driven management.

Virtual Storage Platform G600 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 G family including the Hitachi Unified Storage G600 used in this architecture.

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 NAS Platform 4060

Hitachi NAS Platform is an advanced and integrated network attached storage (NAS) solution. It provides a powerful tool for file sharing, file server consolidation, data protection, and business-critical NAS workloads.

- Powerful hardware-accelerated file system with multi-protocol file services, dynamic provisioning, intelligent tiering, virtualization, and cloud infrastructure
- Seamless integration with Hitachi SAN storage, Hitachi Command Suite, and Hitachi Data Discovery Suite for advanced search and index
- Integration with Hitachi Content Platform for active archiving, regulatory compliance, and large object storage for cloud infrastructure

Hitachi NAS Platform 4060 provides a high-performance and reliable object data storage for the backup solution used with this solution.

Hitachi Dynamic Provisioning

On Hitachi storage systems, Hitachi Dynamic Provisioning provides wide striping and thin provisioning functionalities.

Using Dynamic Provisioning is like using a host-based logical volume manager (LVM), but without incurring host processing overhead. It provides one or more wide-striping pools across many RAID groups. Each pool has one or more dynamic provisioning virtual volume (DP-VOLs) of a logical size you specify of up to 60 TB created against it without allocating any physical space initially.

Deploying Dynamic Provisioning avoids the routine issue of hot spots that occur on logical devices (LDEVs). These occur within individual RAID groups when the host workload exceeds the IOPS or throughput capacity of that RAID group. Dynamic Provisioning distributes the host workload across many RAID groups, which provides a smoothing effect that dramatically reduces hot spots.

This solution uses multiple dedicated dynamic provisioning for different type of Oracle data to avoid intermixing of different type of data I/O which would benefit performance during peak database operations.

Hitachi Dynamic Tiering

Hitachi Dynamic Tiering eliminates manual data classification and movement between storage tiers. This optimizes tiered storage usage while improving performance.

Instead of manually provisioning space from several storage technologies with different performance and cost characteristics, Hitachi Dynamic Tiering enables the management of multiple storage tiers as a single entity. By leveraging the thin provisioning and wide striping features of Hitachi Dynamic Provisioning, Hitachi Dynamic Tiering presents a virtual volume with embedded smart tiering. It monitors access and moves data at the 42MB page level.

Breaking the volume into pages, Hitachi Dynamic Tiering automatically moves infrequently referenced pages to lower cost tiers of storage. Moving pages instead of entire data sets or files reduces the time and storage space required to migrate data.

After an initial setup process, Hitachi Dynamic Tiering monitors data access in real time. It makes decisions on moving data between the available storage tiers based on actual use. Using this approach, Hitachi Dynamic Tiering improves the availability and performance of your storage systems and the applications using that storage.

Hitachi Dynamic Tiering on Hitachi Virtual Storage Platform allows a single pool to contain tiers made up of differently-arranged RAID groups using any type of disk. It manages data migration between the various tiers within a pool automatically. This eliminates most user management of storage tiers within a storage system, and maintains peak performance under dynamic conditions without storage administrator intervention.

The dynamic tiering pool used for this solution provides the capability to later introduce different types of disk to increase the disk pool capacity that is used for the Oracle data without sacrifice of performance since frequently access data will 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

Used for SAN multipathing, the Hitachi Dynamic Link Manager Advanced configuration in this solution uses its round-robin load balancing policy. This policy selects a path by rotating through all available paths. Balancing the load across all available paths optimizes IOPS and response time.

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 using 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 shirt and technologies evolve.

Oracle Database with the Real Application Cluster (RAC) Option

Oracle Database is optimized for use with other Oracle products. It uses Oracle Database Automatic Storage Management (ASM), combining the features of a volume manager and an application-optimized file system for database files. ASM is part of the grid infrastructure component in Oracle Database.

- Real Application Cluster (RAC) scales the database across multiple servers and protects against server failure.
- Automatic Storage Management (ASM) combines the features of a volume manager and an application-optimized file system for database files.

Solution Design

This section describes the reference architecture environment, implementing a quarter-rack environment for Hitachi Unified Compute Platform G6000 for Oracle Real Application Cluster using Hitachi Virtual Storage Platform G600 with Hitachi Accelerated Flash.

Specific infrastructure configuration includes the following:

- **Server** — Two server nodes configured in an Oracle Real Application Cluster with one additional N+M standby node. Another single blade is used as the backup server host.
- **Storage System** — There are LDEVs mapped to each port that are presented to the server as LUNs.
- **SAN Connection** — Each Fiber Channel HBA port is directly connected 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 diagram of this solution.

Figure 2 on page 10 describes the storage configuration for this solution.

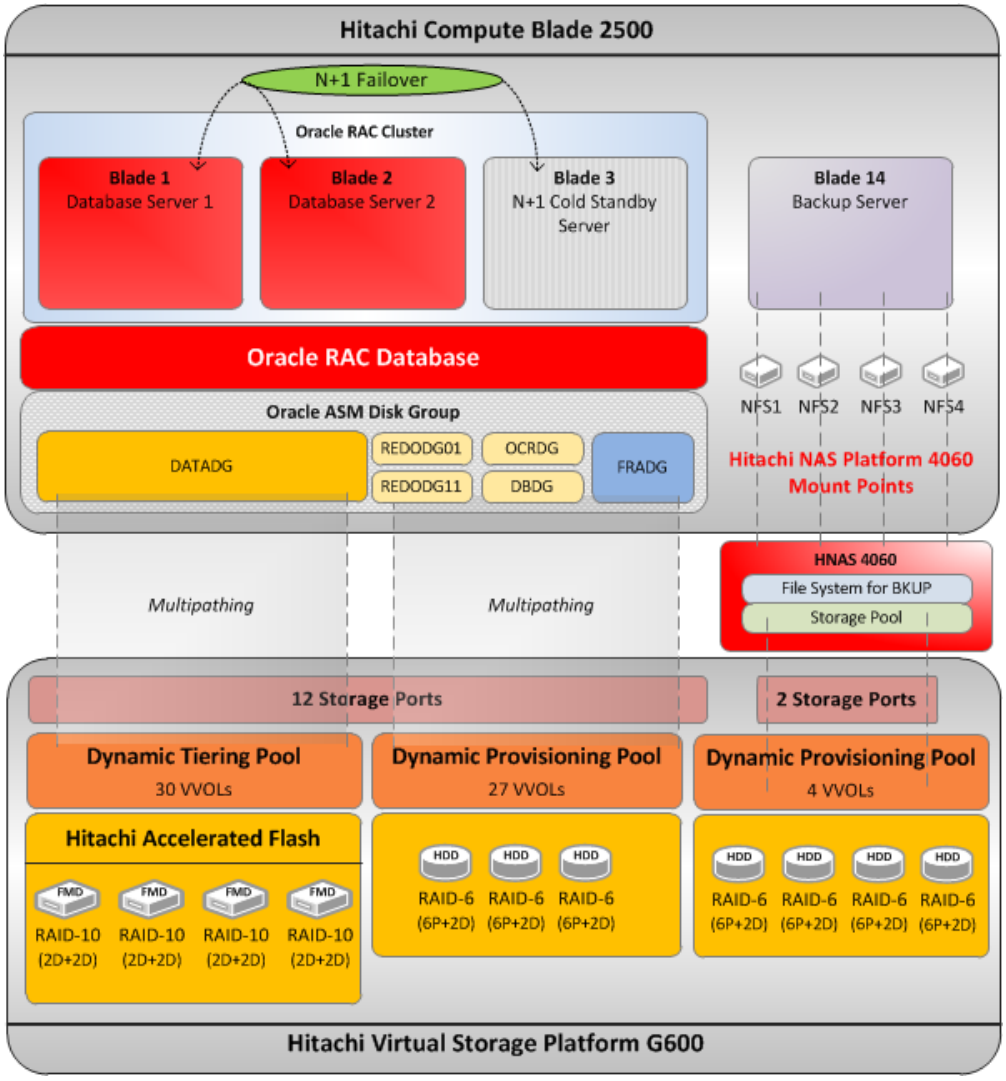
Hitachi Compute Blade 2500
N+1 Failover
Oracle RAC Cluster
Blade 1
Database Server 1
Blade 2
Database Server 2
Blade 3
N+1 Cold Standby Server
Blade 14
Backup Server
Oracle RAC Database
Oracle ASM Disk Group
DATADG
REDODG01
REDODG11
OCRDG
DBDG
FRADG
NFS1
NFS2
NFS3
NFS4
Hitachi NAS Platform 4060
Mount Points
HNAS 4060
File System for BKUP
Storage Pool
Multipathing
Multipathing
12 Storage Ports
2 Storage Ports
Dynamic Tiering Pool
30 VVOLs
Dynamic Provisioning Pool
27 VVOLs
Dynamic Provisioning Pool
4 VVOLs
Hitachi Accelerated Flash
FMD
RAID-10 (2D+2D)
HDD
RAID-6 (6P+2D)
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	hnas_sas_pool_03
Pool Type	Dynamic Tiering	Dynamic Provisioning	Dynamic Provisioning
RAID Group	1-1 – 1-4	1-5 – 1-7	1-8 - 1-11
RAID Level	RAID-10 (2D+2D)	RAID-6 (6P+2D)	RAID-6 (6P+2D)
Number of LDEVs	4	6	4
Drive Type	1.6 TB Flash Module Drive (FMD)	1.2 TB 10k RPM SAS Drive	1.2 TB 10k RPM SAS Drive
Number of Drives	16, 4 per each tray	24	32
Pool Capacity	11.99 TB	17.93 TB	38 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	hnas_sas_pool_03
Number of VVOLs	30	27	4
VVOL Size	400 GB	200 GB	9 TB
Purpose	OLTP Application Tablespaces	■ Oracle System ■ Sysaux ■ Undo ■ Temp ■ Redo Logs ■ Oracle Cluster Registry ■ Voting Disk	■ System drive for Hitachi NAS Platform.
Storage Port	1A, 1B, 3A, 3B, 2A, 2B, 4A, 4B		1C, 2C

Hitachi NAS Platform 4060 Configuration

Hitachi NAS Platform 4060 provides a highly reliable file storage system for the Oracle database backup. It consists of thirty two 1.2 TB SAS 10k RPM drives presented from Hitachi Virtual Storage Platform G600 to the Hitachi NAS Platform 4060 head with the following configuration:

- **Node 1 Pool 0** — Four LDEVs from a dynamic provisioning pool consisting of four RAID groups in a RAID-6 (6D+2P) configuration presented as four Hitachi NAS Platform system drives

The NAS Platform 4060 head contains one storage pool, one file system, and one enterprise virtual server (EVS). The EVS hosts these file systems. When requiring additional capacity, add additional system drives to the storage pool and file system capacity on the fly.

The recommended file system block size is 32 KB for backup uses. NAS Platform virtual volume quotas should be set to 85% full.

Hitachi Data Systems best practice for a production database backup recommends the use Hitachi ShadowImage Heterogeneous Replication to copy data onto a separate storage pool consisting of nearline SAS disks. The backup operation can be done from replicated data to minimize backup operation overhead to the production database.

Consult with Hitachi Data Systems Global Services Solutions for sizing details for ShadowImage volumes.

Table 5 shows the NFS exports for the Oracle Database backup.

Table 5. NFS Exports for the Oracle Database Backup

LDEVs	System Drives	Storage Pool	File System	NFS Exports (Name/Path)	NAS Platform Head	Purpose
4	1 to 4	BKUP_SP	fsbkup	exp/bkup_01 - exp/bkup_04	HNAS-1	Oracle database backup

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 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. Limit the size of each small file tempfile to 31.25 GB.
- **Undo Tablespace** — Place UNDO tablespace in this configuration in the Data ASM diskgroup. Assign one 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 Archive 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 follows:
  - FILESYSTEMIO_OPTIONS = setall

Table 6 shows the Oracle RAC database configuration.

Table 6. Oracle RAC Database Settings

Environment	Value
RAC configuration	Yes
ASM	Yes - Oracle RAC Database
HNAS NFS	Yes - Oracle Database Backup using RMAN

Table 7 shows the Oracle Environment Parameters.

Table 7. Oracle Environment Parameters

Setting	Parameter
SGA_TARGET	69.75 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	274874368
USE_LARGE_PAGES	TRUE
FILESYSTEMIO_OPTIONS	SETALL

Table 8 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

ASM Disk Group	ASM Disk	HDLM LUNs	LUNs Count	Purpose
OCRDG	OCDISK1	/dev/sddlmbp1	1	Oracle Cluster Registry and Voting Disk
REDODG_01	RGDISK01	/dev/sddlmca1	1	Online REDO log group and control file
REDODG_11	RGDISK11	/dev/sddlmcb1	1	

Table 8. Oracle ASM Disk Configuration (Continued)

ASM Disk Group	ASM Disk	HDLM LUNs	LUNs Count	Purpose
FRADG	FRDISK1	/dev/sddlmcc1	15	Flash Recovery Area
	FRDISK2	/dev/sddlmcd1		
	FRDISK3	/dev/sddlmce1		
	FRDISK4	/dev/sddlmcf1		
	FRDISK5	/dev/sddlmcg1		
	FRDISK6	/dev/sddlmch1		
	FRDISK7	/dev/sddlmci1		
	FRDISK8	/dev/sddlmcj1		
	FRDISK9	/dev/sddlmck1		
	FRDISK10	/dev/sddlmcl1		
	FRDISK11	/dev/sddlmcm1		
	FRDISK12	/dev/sddlmcn1		
	FRDISK13	/dev/sddlmco1		
	FRDISK14	/dev/sddlmcp1		
	FRDISK15	/dev/sddlmda1		
DBDG	SYSDISK1	/dev/sddlmdb1	6	Sys, Undo and Temp
	UNDISK1	/dev/sddlmdc1		
	UNDISK2	/dev/sddlmdd1		
	TMPDISK1	/dev/sddlmde1		
	TMPDISK2	/dev/sddlmdf1		
	TMPDISK3	/dev/sddlmdg1		

Table 8. Oracle ASM Disk Configuration (Continued)

ASM Disk Group	ASM Disk	HDLM LUNs	LUNs Count	Purpose
BWDG	DADISK1	/dev/sddlmaa1	30	Application Data
	DADISK2	/dev/sddlmab1		
	DADISK3	/dev/sddlmac1		
	DADISK4	/dev/sddlmad1		
	DADISK5	/dev/sddlmae1		
	DADISK6	/dev/sddlmaf1		
	DADISK7	/dev/sddlmag1		
	DADISK8	/dev/sddlmah1		
	DADISK9	/dev/sddlmai1		
	DADISK10	/dev/sddlmaj1		
	DADISK11	/dev/sddlmak1		
	DADISK12	/dev/sddlmal1		
	DADISK13	/dev/sddlmam1		
	DADISK14	/dev/sddlman1		
	DADISK15	/dev/sddlmao1		
	DADISK16	/dev/sddlmap1		
	DADISK17	/dev/sddlmba1		
	DADISK18	/dev/sddlmbb1		
	DADISK19	/dev/sddlmbc1		
	DADISK20	/dev/sddlmbd1		
	DADISK21	/dev/sddlmbe1		
	DADISK22	/dev/sddlmbf1		
	DADISK23	/dev/sddlmbg1		
	DADISK24	/dev/sddlmbh1		
	DADISK25	/dev/sddlmbi1		
	DADISK26	/dev/sddlmbj1		
	DADISK27	/dev/sddlmbk1		
	DADISK28	/dev/sddlmbl1		
	DADISK29	/dev/sddlmbm1		
	DADISK30	/dev/sddlmbo1		

Server and Application Architecture

This reference architecture uses a single Hitachi Compute Blade 2500 chassis with four server blades. Two server blades are used for a two-node RAC configuration with an additional N+M cold standby node. The fourth server blade is used for the backup server.

This provides the compute power for Oracle RAC database to handle complex database queries and a large volume of transaction processing in parallel. Table 9 on page 16 describes the details of the server configuration for this solution.

Table 9. Server Details

Server	Form Size	Server Name	Role	CPU Core	RAM	Blade
Node1	Half-Width size	RAC01	RAC Node	36	256 GB	1
Node2	Half-Width size	RAC02	RAC Node	36	256 GB	2
Node3	Half-Width size	N+M	Standby	36	256 GB	3
Node4	Half-Width size	Backup01	Database Backup	36	256 GB	14

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 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 prior to configuring of 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 3 shows the server infrastructure for the reference architecture with the cold standby node. The recommended slot for the backup server is Server Blade 14 to allow ease of expansion to Oracle RAC at a later time.

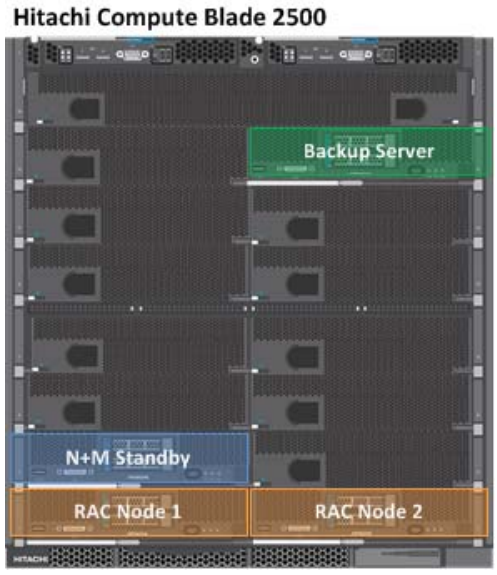

Figure 3

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 on page 4. This provides a four path connection for all LUNs mapped to each of the database servers in Oracle RAC database. Table 10 shows the direct-connect from the HBA of the server blade to the Hitachi Virtual Storage Platform G600 ports.

Table 10. 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_B0_1A
		HBA1-2	2A	BS25K_B0_2A
		HBA2-1	1B	BS25K_B0_1B
		HBA2-2	2B	BS25K_B0_2B
	BLADE 2	HBA1-1	3A	BS25K_B1_3A
		HBA1-2	4A	BS25K_B1_4A
		HBA2-1	3B	BS25K_B1_3B
		HBA2-2	4B	BS25K_B1_4B
	BLADE 3	HBA1-1	1C	BS25K_B0_1A BS25K_B1_3A
		HBA1-2	2C	BS25K_B0_2A BS25K_B1_4A
		HBA2-1	1D	BS25K_B0_1B BS25K_B1_3B
		HBA2-2	2D	BS25K_B0_2B BS25K_B1_4B
	BLADE 14	HBA1-1	3C	BS25K_B14_3C
		HBA2-1	4C	BS25K_B14_4C
Hitachi NAS Platform 4060	EVSHNAS01	8GFC-1	3D	HNAS01_FC1_3D
		8GFC-2	4D	HNAS01_FC2_4D

Hitachi 16 Gb/sec PCIe HBA Card Configuration

This section 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 prior to configuring of 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 on page 18 shows the Hitachi 16 Gb/sec HBA PCIe cards that are installed in the Hitachi Compute Blade 2500 chassis.

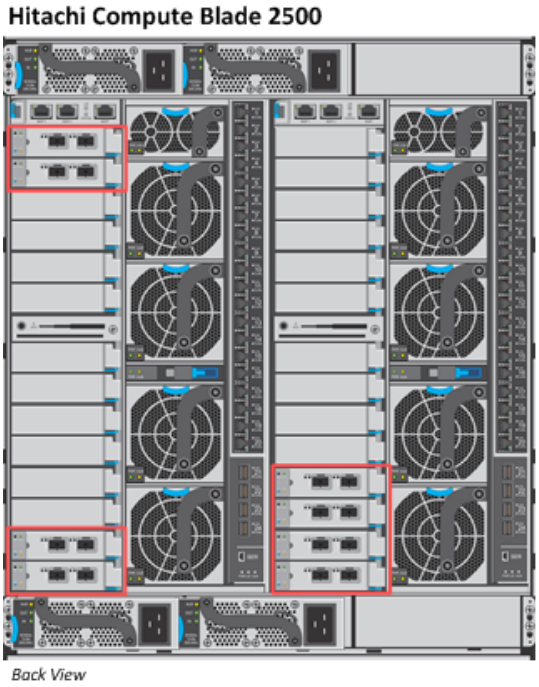

Figure 4

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

Table 11. Hitachi HBA PCIe Card Parameters

Setting	Parameter
Boot Function	Enable
Link Speed	16Gbps
Connection Type	Point-to-Pont
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.
- **Backup Network** — This network provides a dedicated logical network for backup purposes. This ensures that backup data is not being transmitted over the network used for application data.

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.

The Hitachi NAS Platform 4060 Ethernet ports are connected to each of the external ports of the 10 Gb/sec switches.

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 on page 20 shows the network configuration for the reference architecture environment.

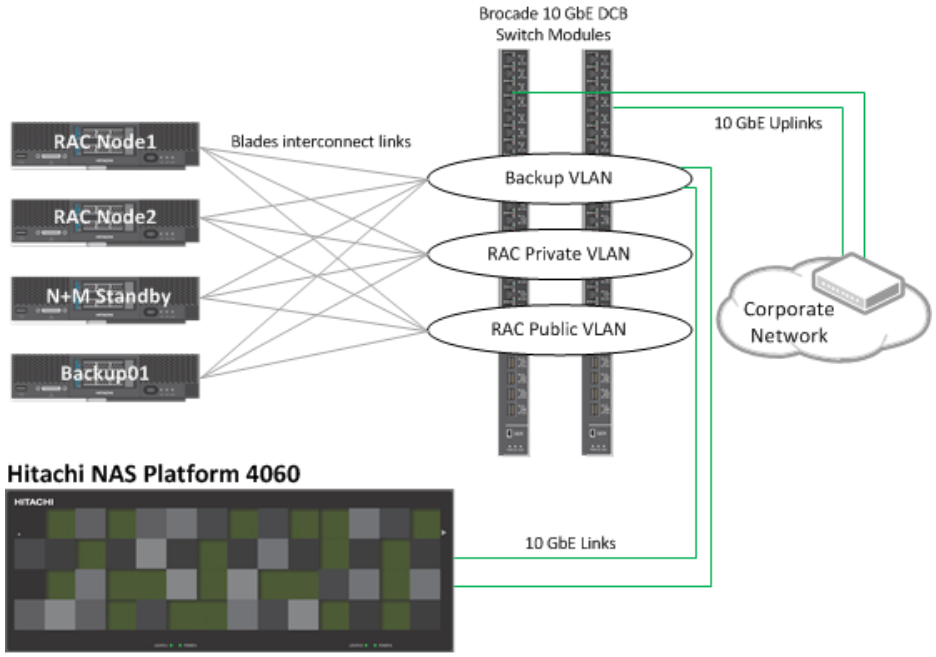

Figure 5

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

Table 12. 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	15	243	Bond200	Public Oracle	9
		6	1	15	244	Bond300	Public Management	1
	B1-CNIC-3	3	2	15	243	Bond200	Public Oracle	9
		7	2	15	244	Bond300	Public Management	1
Database Server 2	B2-CNIC-0	0	1	2	208	Bond1	Private	10
	B2-CNIC-1	1	2	2		Available		10
	B2-CNIC-2	2	1	16	243	Bond200	Public	9
		6	1	16		Bond300		1
	B2-CNIC-3	3	2	16		Bond200		9
		7	2	16		Bond300		1
N+1 Cold Standby Server	B2-CNIC-0	0	1	3	208	From failed server	Private	10
	B2-CNIC-1	1	2	3				10
	B2-CNIC-2	2	1	17			Public	9
		6	1	17				1
	B2-CNIC-3	3	2	17				9
		7	2	17				1
Backup Server	B2-CNIC-0	0	1	4	500	Bond1	Private Backup	10
	B2-CNIC-1	1	2	4				
	B2-CNIC-2	2	1	18	243	Bond200	Public Oracle	9
		6	1	18		Bond300-	Public Management	1
	B2-CNIC-3	3	2	18		Bond200-	Public Oracle	9
		7	2	18		Bond300-	Public Management	1

Table 12. Network Configuration (Continued)

Server	NIC Ports	UMC PF#	Switch Bay ID	Switch Ports (Internal)	VLAN/ Subnet	NIC BOND	Network	Bandwidth (Gb/sec)
Hitachi NAS Platform 4060	HNAS1-NIC-0	Not applicable			500	Not applicable	Private (Admin)	1
	SMU300-NIC-0				500			1
	SMU400-NIC-0				500			1

Test Results

This section summarizes the key observations from the test results for Hitachi Unified Compute Platform G6000 for Oracle Real Application Cluster using Hitachi Virtual Storage Platform G600 with Hitachi Accelerated Flash and Oracle Orion.

Table 13 lists the Oracle Orion I/O test cases for this solution.

Table 13. Oracle Orion I/O Test Case

Test Case	Metric	Value
100% Sequential Reads	Throughput in MB/sec	12,526
90% Sequential Reads	Throughput in MB/sec	12,573
100% Random Reads	IOPS	413,818
100% Random Reads	Latency in milliseconds	2.45

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-410-01, September 2015.