

Hitachi Unified Compute Platform 6000 for Oracle Real Application Clusters on Four Nodes Using Hitachi Virtual Storage Platform G800, Hitachi Accelerated Flash, and Hitachi Compute Blade 2500 With Intel Xeon E7-8893 v4

Reference Architecture Guide

By Eduardo Freitas, Libo Jiang, and Steve Burns

September 2016

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

Solution Overview	3
Key Solution Components.....	5
Hitachi Compute Blade 2500	6
Hitachi Virtual Storage Platform Gx00 Models	6
Hitachi Accelerated Flash	7
Hitachi Storage Virtualization Operating System	7
Hitachi Dynamic Link Manager Advanced	7
Hitachi Data Instance Director	8
Brocade Switches	8
Red Hat Enterprise Linux	8
Oracle Database	8
Solution Design.....	9
Storage Architecture	9
Database Layout	11
Database Storage Snapshots	17
Server and Application Architecture	17
SAN Architecture.....	18
Network Architecture	20

Hitachi Unified Compute Platform 6000 for Oracle Real Application Clusters on Four Nodes Using Hitachi Virtual Storage Platform G800, Hitachi Accelerated Flash, and Hitachi Compute Blade 2500 With Intel Xeon E7-8893 v4

Reference Architecture Guide

This reference architecture guide shows how using Hitachi Unified Compute Platform for Oracle Real Application Clusters (UCP for Oracle RAC) in a 4-node solution provides a high performance, integrated solution. The environment uses Hitachi Virtual Storage Platform G800 (VSP G800) with Hitachi Accelerated Flash (HAF). Use this document to design an infrastructure for your requirements and budget.

This 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 workloads, such as OLTP and DSS. The dedicated servers run Oracle Database 12c R1 with the Oracle Real Application Clusters (RAC) option. The operating system is Red Hat Enterprise Linux 7.2.

Hitachi Data Instance Director (HDID) is configured with one master node and one repository node on Microsoft® Windows Server® 2012 R2. Data Instance Director can be installed on either a physical server or Microsoft Hyper-V® virtual machines.

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
- Backup administrators
- IT professionals with the responsibility of backing up, restoring and recovering, or disaster database recovery of an Oracle Database solution

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

- Hitachi Virtual Storage Platform G800 (VSP G800) using Hitachi Accelerated Flash (HAF)
- Hitachi Compute Blade 2500 (CB 2500) using 520X B3 server blades with Intel Xeon E7-8893 v4 processor CPUs
- Hitachi Data Instance Director (HDID)
- Hitachi Thin Image (HTI) - used for replication
- 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 (UCP for Oracle RAC) on two nodes using Hitachi Virtual Storage Platform G800 (VSP G800) with Hitachi Accelerated Flash (HAF). In this configuration, Hitachi Compute Blade 2500 (CB 2500) uses 520X B3 server blades with Intel Xeon E7-8893 v4 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** with five server blades
 - **Server Blade 1** — Oracle RAC NODE 1
 - **Server Blade 3** — Oracle RAC NODE 2
 - **Server Blade 5** — Oracle RAC NODE 3
 - **Server Blade 7** — Oracle RAC NODE 4
 - **Server Blade 9** — N+M cold standby node
- **Hitachi Virtual Storage Platform G800** with Hitachi Accelerated Flash
- 16 Gb/sec SAN infrastructure
- 10 GbE LAN infrastructure

Figure 1 on page 4 shows the high-level infrastructure for this solution.




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 Systems

Hardware	Detail Description	Version	Quantity
Hitachi Virtual Storage Platform G800 (VSP G800)	<ul style="list-style-type: none"> ▪ Dual controller ▪ 20 × 16 Gb/sec Fibre Channel ports ▪ 8 × 12 Gb/sec backend serial attached SCSI (SAS) ports ▪ 512 GB cache memory 	83-03-24-00/00	1
	<ul style="list-style-type: none"> ▪ 3.2 TB flash memory drives (FMDs) <ul style="list-style-type: none"> ▪ 2 spares 		34 — 9 on the first two trays, and 8 on the last two
	<ul style="list-style-type: none"> ▪ 1.2 TB 10K RPM SAS drives <ul style="list-style-type: none"> ▪ 2 spares 		92
Hitachi Compute Blade 2500 chassis (CB 2500)	<ul style="list-style-type: none"> ▪ 2 × 10 GbE data center bridging (DCB) local area network (LAN) switch module Brocade VDX2746 ▪ 10 fan modules ▪ 2 management modules ▪ 4 power supply modules 	Management Module Firmware Version A0160-B-1453 Dictionary Version A0031 DCB Switch Version 6.0.2b	1
520X B3 full-width server blade	<ul style="list-style-type: none"> ▪ 2 Intel Xeon E7-8893 v4 processor CPUs ▪ 256 GB RAM, 16 × 16 GB DDR4 memory ▪ 1 × 4-port 10 GbE converged network adapter (CNA) LAN on motherboard (LOM) 	11-01/11-02	5
	<ul style="list-style-type: none"> ▪ Emulex LPe16002B 16 Gb/sec 2-port PCIe Fibre Channel HBA 	Firmware: 10.7.0.1	10
Brocade 6510 Fibre Channel switch	<ul style="list-style-type: none"> ▪ 48 port Fibre Channel switch ▪ 20 × 16 Gb/sec SFPs on first switch ▪ 21 × 16 Gb/sec SFPs on second switch 	7.4.1b	2

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 (SVOS) with the following: <ul style="list-style-type: none"> ■ Hitachi Storage Navigator (SN) ■ Hitachi Dynamic Tiering (HDT) ■ Hitachi Dynamic Provisioning (HDP) 	Microcode dependent	Storage management suite
Hitachi Data Instance Director (HDID)	5.3	Data protection
Hitachi Command Suite (HCS)	8.4.1-02	Storage management suite
Hitachi Compute Systems Manager (HCSM)	8.4.1-02	N+M Management
Red Hat Enterprise Linux (RHEL)	7.2	Operating system for Oracle RAC
Oracle Grid Infrastructure	12c Release 1 12.1.0.2.0	Oracle Clusterware
Oracle Database	12c Release 1 12.1.0.2.0	Oracle database system
Hitachi Dynamic Link Manager Advanced (HDLMA)	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 uses 520X B3 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 G800, 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 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 types 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 the capability to introduce different types of disks at when it is 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 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 by 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.

Hitachi Data Instance Director

[Hitachi Data Instance Director](#) provides a modern, holistic approach to data protection, recovery, and retention.

Data Instance Director has a unique workflow-based policy engine, presented in a whiteboard-style user interface. It helps you to map copy data management processes to business priorities.

Data Instance Director includes a wide range of fully integrated storage-based and host-based incremental-forever data capture capabilities. These can be combined into complex workflows to automate and simplify copy data management.

Use Data Instance Director with Oracle so you no longer need to be dependent on IT personnel to backup, restore, recover, or clone databases. This allows database administrators and storage administrators to focus on their specialty while unburdening IT staff of time-consuming tasks.

Data Instance Director provides application-consistent storage-based clone and remote file replication orchestration for Oracle Database environments using an intuitive, comprehensive, and efficient user interface.

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 VDX 2746 10 GbE switch module
- Brocade 6510 48 port Fibre Channel 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 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.

[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 describes the reference architecture environment, implementing a half-rack environment for Hitachi Unified Compute Platform 6000 for Oracle Real Application Clusters using Hitachi Virtual Storage Platform G800 with Hitachi Accelerated Flash.

Specific infrastructure configuration details include the following:

- **Server** — Configure four server nodes in an Oracle Real Application Clusters 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 to the storage front-end ports using the switched SAN environment.

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 Configuration

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

Figure 2 on page 10 shows the layout of the storage configuration used for this solution.

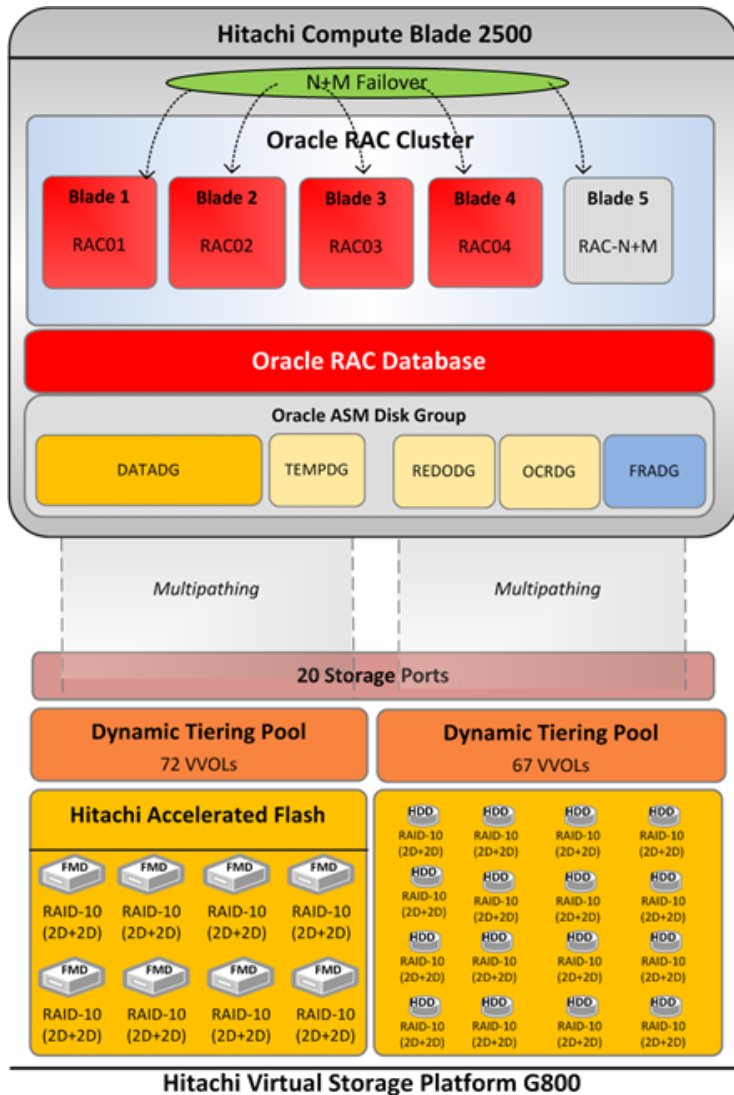


Figure 2

Table 3 shows the storage pool configuration used in the tested configuration. You may use a different configuration.

Table 3. Storage Pool Configuration

Pool ID	ora_prod_fmd_01	ora_prod_sas_01	ora_bkup_sas_01
Pool Type	Dynamic Tiering	Dynamic Tiering	Dynamic Provisioning
RAID Group	1-1 – 1-8	1-9 – 1-24	1-25 - 1-28
RAID Level	RAID-10 (2D+2D)	RAID-10 (2D+2D)	RAID-6 (6D+2P)
Number of LDEVs	256 × 204.7 GB	64 × 536.8 GB	16 × 1610.4 GB
Drive Type	3.2 TB Flash Module Drive (FMD)	1.2 TB 10K RPM SAS Drive	1.2 TB 10K RPM SAS Drive
Number of Drives	32 with 8 per tray	64 with 16 per tray	32 with 8 per tray
Pool Capacity	51.17 TB	33.55 TB	25.16 TB

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

Table 4. Logical Storage Configuration

Pool ID	ora_prod_sas_01	ora_prod_fmd_01	ora_prod_fmd_01	ora_prod_sas_01	ora_prod_sas_01	ora_prod_sas_01	ora_prod_sas_01
Number of VVOLs	3	64	8	32	32	4	1
VVOL Size	5 GB	600 GB	200 GB	8 GB	960 GB	200 GB	100 MB
Purpose	Oracle Cluster Registry Voting Disk	OLTP Application Tablespaces System Sysaux Undo	Temp	Online Redo Logs Control Files	Incremental Backups Archived Redo Logs Control File Autobackups	SAN Boot	CMD Device for the command control interface
Storage Port	20 ports for Oracle RAC						1 port for HDID

Database Layout

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

- **Data and Indexes Tablespace** — Assign an ASM diskgroup for the data and index tablespaces. The smallfile tablespace for data consists of 2048 data files that are 8 GB each. Set the tablespace to a small initial size with **auto extend** enabled to maximize storage utilization.
- **TEMP Tablespace** — Place the TEMP tablespace of this configuration in the TEMP ASM diskgroup. A number of tempfiles are created within one smallfile 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. A number of undo datafiles are created within each smallfile UNDO tablespace. Limit the size of each undo datafile to 8 GB.
- **Online Redo Logs** — Assign one ASM diskgroup for each database. Four redo logs are created for each database instance in a four 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 OCR 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 5 has the Oracle RAC database configuration.

Table 5. Oracle RAC Database Settings

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

Table 6 lists the Oracle environment parameters.

Table 6. 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
FILESYSTEMIO_OPTIONS	SETALL

Table 7 has the details of 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/sddlmac1	3	<ul style="list-style-type: none"> ■ Oracle Cluster Registry and Voting Disk
	OCRDISK2	/dev/sddlmac1		
	OCRDISK3	/dev/sddlmac1		

Table 7. Oracle ASM Disk Configuration (Continued)

ASM Disk Group	ASM Disk	HDLM LUNs	LUNs Count	Purpose
REDODG	RGDISK01	/dev/sddlmae1	32	<ul style="list-style-type: none"> ■ Online REDO Logs ■ Control Files
	RGDISK02	/dev/sddlmaf1		
	RGDISK03	/dev/sddlmag1		
	RGDISK04	/dev/sddlmah1		
	RGDISK05	/dev/sddlmai1		
	RGDISK06	/dev/sddlmaj1		
	RGDISK07	/dev/sddlmak1		
	RGDISK08	/dev/sddlmal1		
	RGDISK09	/dev/sddlmam1		
	RGDISK10	/dev/sddlman1		
	RGDISK11	/dev/sddlmao1		
	RGDISK12	/dev/sddlmap1		
	RGDISK13	/dev/sddlmba1		
	RGDISK14	/dev/sddlmbb1		
	RGDISK15	/dev/sddlmbc1		
	RGDISK16	/dev/sddlmbd1		
	RGDISK17	/dev/sddlmbel1		
	RGDISK18	/dev/sddlmbf1		
	RGDISK19	/dev/sddlmbg1		
	RGDISK20	/dev/sddlmbh1		
	RGDISK21	/dev/sddlmbi1		
	RGDISK22	/dev/sddlmbj1		
	RGDISK23	/dev/sddlmbk1		
	RGDISK24	/dev/sddlmb1		
	RGDISK25	/dev/sddlmbm1		
	RGDISK26	/dev/sddlmbn1		
	RGDISK27	/dev/sddlmbol1		
	RGDISK28	/dev/sddlmbp1		
	RGDISK29	/dev/sddlmbca1		
	RGDISK30	/dev/sddlmbc1		
	RGDISK31	/dev/sddlmbcc1		
	RGDISK32	/dev/sddlmbcd1		

Table 7. Oracle ASM Disk Configuration (Continued)

ASM Disk Group	ASM Disk	HDLM LUNs	LUNs Count	Purpose
DATADG	DADISK01	/dev/sddlmc1	64	<ul style="list-style-type: none"> ■ Application Data ■ Undo ■ System ■ Sysaux
	DADISK02	/dev/sddlmc2		
	DADISK03	/dev/sddlmc3		
	DADISK04	/dev/sddlmc4		
	DADISK05	/dev/sddlmc5		
	DADISK06	/dev/sddlmc6		
	DADISK07	/dev/sddlmc7		
	DADISK08	/dev/sddlmc8		
	DADISK09	/dev/sddlmc9		
	DADISK10	/dev/sddlmc10		
	DADISK11	/dev/sddlmc11		
	DADISK12	/dev/sddlmc12		
	DADISK13	/dev/sddlmc13		
	DADISK14	/dev/sddlmc14		
	DADISK15	/dev/sddlmc15		
	DADISK16	/dev/sddlmc16		
	DADISK17	/dev/sddlmc17		
	DADISK18	/dev/sddlmc18		
	DADISK19	/dev/sddlmc19		
	DADISK20	/dev/sddlmc20		
	DADISK21	/dev/sddlmc21		
	DADISK22	/dev/sddlmc22		
	DADISK23	/dev/sddlmc23		
	DADISK24	/dev/sddlmc24		
	DADISK25	/dev/sddlmc25		
	DADISK26	/dev/sddlmc26		
	DADISK27	/dev/sddlmc27		
	DADISK28	/dev/sddlmc28		
	DADISK29	/dev/sddlmc29		
	DADISK30	/dev/sddlmc30		
	DADISK31	/dev/sddlmc31		
	DADISK32	/dev/sddlmc32		

Table 7. Oracle ASM Disk Configuration (Continued)

ASM Disk Group	ASM Disk	HDLM LUNs	LUNs Count	Purpose
DATADG (continued)	DADISK33	/dev/sddlme1		
	DADISK34	/dev/sddlme2		
	DADISK35	/dev/sddlme3		
	DADISK36	/dev/sddlme4		
	DADISK37	/dev/sddlme5		
	DADISK38	/dev/sddlme6		
	DADISK39	/dev/sddlme7		
	DADISK40	/dev/sddlme8		
	DADISK41	/dev/sddlme9		
	DADISK42	/dev/sddlme10		
	DADISK43	/dev/sddlme11		
	DADISK44	/dev/sddlme12		
	DADISK45	/dev/sddlme13		
	DADISK46	/dev/sddlme14		
	DADISK47	/dev/sddlme15		
	DADISK48	/dev/sddlme16		
	DADISK49	/dev/sddlme17		
	DADISK50	/dev/sddlme18		
	DADISK51	/dev/sddlme19		
	DADISK52	/dev/sddlme20		
	DADISK53	/dev/sddlme21		
	DADISK54	/dev/sddlme22		
	DADISK55	/dev/sddlme23		
	DADISK56	/dev/sddlme24		
	DADISK57	/dev/sddlme25		
	DADISK58	/dev/sddlme26		
	DADISK59	/dev/sddlme27		
	DADISK60	/dev/sddlme28		
	DADISK61	/dev/sddlme29		
	DADISK62	/dev/sddlme30		
	DADISK63	/dev/sddlme31		
	DADISK64	/dev/sddlme32		

Table 7. Oracle ASM Disk Configuration (Continued)

ASM Disk Group	ASM Disk	HDLM LUNs	LUNs Count	Purpose
TEMPDG	TEDISK01	/dev/sddlmg1	8	<ul style="list-style-type: none"> ■ TEMP
	TEDISK02	/dev/sddlmgf1		
	TEDISK03	/dev/sddlmgg1		
	TEDISK04	/dev/sddlmg1		
	TEDISK05	/dev/sddlmgj1		
	TEDISK06	/dev/sddlmgk1		
	TEDISK07	/dev/sddlmgk1		
	TEDISK08	/dev/sddlmg1		
FRADG	FRDISK01	/dev/sddlmgm1	32	<ul style="list-style-type: none"> ■ Archive Logs ■ Incremental Backups ■ Control File Autobackups
	FRDISK02	/dev/sddlmg1		
	FRDISK03	/dev/sddlngo1		
	FRDISK04	/dev/sddlmgp1		
	FRDISK05	/dev/sddlmha1		
	FRDISK06	/dev/sddlmh1		
	FRDISK07	/dev/sddlmh1		
	FRDISK08	/dev/sddlmh1		
	FRDISK09	/dev/sddlmh1		
	FRDISK10	/dev/sddlmh1		
	FRDISK11	/dev/sddlmh1		
	FRDISK12	/dev/sddlmh1		
	FRDISK13	/dev/sddlmh1		
	FRDISK14	/dev/sddlmh1		
	FRDISK15	/dev/sddlmh1		
	FRDISK16	/dev/sddlmh1		
	FRDISK17	/dev/sddlmh1		
	FRDISK18	/dev/sddlmh1		
	FRDISK19	/dev/sddlmh1		
	FRDISK20	/dev/sddlmh1		
	FRDISK21	/dev/sddlmi1		
	FRDISK22	/dev/sddlmi1		
	FRDISK23	/dev/sddlmi1		
	FRDISK24	/dev/sddlmi1		
	FRDISK25	/dev/sddlmi1		
	FRDISK26	/dev/sddlmi1		
	FRDISK27	/dev/sddlmi1		

Table 7. Oracle ASM Disk Configuration (Continued)

ASM Disk Group	ASM Disk	HDLM LUNs	LUNs Count	Purpose
FRADG (continued)	FRDISK28	/dev/sddlmi1		
	FRDISK29	/dev/sddlmi2		
	FRDISK30	/dev/sddlmi3		
	FRDISK31	/dev/sddlmi4		
	FRDISK32	/dev/sddlmi5		

Database Storage Snapshots

A storage snapshot is a backup copy created at a particular point in time. Snapshots do not require an initial copy. They are not stored as physical copies of blocks, but rather as pointers to the blocks that existed when creating the snapshot. Because of this tight physical relationship, the snapshot is maintained on the same storage array as the original data.

Hitachi Data Instance Director (HDID) uses Hitachi snapshot technology with an Oracle database to create fast and space-efficient backups with no data loss or data integrity issues. The database snapshot is available in non-ASM, ASM, and RAC configurations.

Use the snapshot operation to trigger snapshots of entire volumes.

You can access snapshots using the Restore screen. The snapshots can be mounted or used to revert a volume to a previous state.

Concurrent Oracle database snapshots keep identical backups, which simplifies recovering the operational Oracle database during disaster recovery.

Server and Application Architecture

This reference architecture uses a single Hitachi Compute Blade 2500 chassis with five server blades. Four server blades are used for a four-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 lists the details of the server configuration for this solution.

Table 8. Server Details

Blade	Form Size	Server Name	Role	CPU Core	RAM
1	Full-Width	RAC01	RAC Node	8	256 GB
3	Full-Width	RAC02	RAC Node	8	256 GB
5	Full-Width	RAC03	RAC Node	8	256 GB
7	Full-Width	RAC04	RAC Node	8	256 GB
9	Full-Width	RAC-N+M	Standby	8	256 GB

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 9 is configured as the cold standby node for the four active nodes.

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 G800 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 Emulex listed in Table 1, “Key Solution Components From Hitachi Data Systems,” on page 5. 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 SAN connections from the HBA of the server blade to the Hitachi Virtual Storage Platform G800 ports.

Table 9. Fibre Channel Configuration on Hitachi Virtual Storage Platform G800

Host	HBA	Storage Port	Storage Host Group	Zone Name
RAC01	HBA1-1	1A	CB2500_B1_1A	RAC01_HBA1_1_G800_1A
	HBA1-2	2A	CB2500_B1_2A	RAC01_HBA1_2_G800_2A
	HBA2-1	1B	CB2500_B1_1B	RAC01_HBA2_1_G800_1B
	HBA2-2	2B	CB2500_B1_2B	RAC01_HBA2_2_G800_2B

Table 9. Fibre Channel Configuration on Hitachi Virtual Storage Platform G800 (Continued)

Host	HBA	Storage Port	Storage Host Group	Zone Name
RAC02	HBA1-1	3A	CB2500_B3_3A	RAC02_HBA1_1_G800_3A
	HBA1-2	4A	CB2500_B3_4A	RAC02_HBA1_2_G800_4A
	HBA2-1	3B	CB2500_B3_3B	RAC02_HBA2_1_G800_3B
	HBA2-2	4B	CB2500_B3_4B	RAC02_HBA2_2_G800_4B
RAC03	HBA1-1	1C	CB2500_B5_1C	RAC03_HBA1_1_G800_1C
	HBA1-2	2C	CB2500_B5_2C	RAC03_HBA1_2_G800_2C
	HBA2-1	1D	CB2500_B5_1D	RAC03_HBA2_1_G800_1D
	HBA2-2	2D	CB2500_B5_2D	RAC03_HBA2_2_G800_2D
RAC04	HBA1-1	3C	CB2500_B7_3C	RAC04_HBA1_1_G800_3C
	HBA1-2	4C	CB2500_B7_4C	RAC04_HBA1_2_G800_4C
	HBA2-1	3D	CB2500_B7_3D	RAC04_HBA2_1_G800_3D
	HBA2-2	4D	CB2500_B7_4D	RAC04_HBA2_2_G800_4D
RAC-N+M	HBA1-1	1E	CB2500_B1_1A	RACNM_HBA1_1_G800_1E
			CB2500_B3_3A	
			CB2500_B5_1C	
			CB2500_B7_3C	
	HBA1-2	2E	CB2500_B1_2A	RACNM_HBA1_2_G800_2E
			CB2500_B3_4A	
			CB2500_B5_2C	
			CB2500_B7_4C	
	HBA2-1	1F	CB2500_B1_1B	RACNM_HBA2_1_G800_1F
			CB2500_B3_3B	
			CB2500_B5_1D	
			CB2500_B7_3D	
	HBA2-2	2F	CB2500_B1_2B	RACNM_HBA2_2_G800_2F
			CB2500_B3_4B	
			CB2500_B5_2D	
			CB2500_B7_4D	

Emulex 16 Gb/sec PCIe HBA Card Configuration

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

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




Figure 4

Set the following parameters for each of the Emulex HBA PCIe cards following Table 10.

Table 10. Emulex HBA PCIe Card Parameters

For This	Use This
Boot Function	Enable
Link Speed	16 Gb/s link speed
Connection Type	Fabric Point to Point

Network Architecture

This architecture requires the following separate networks:

- **Private Network (also called cluster interconnect)** — This network must be scalable. In addition, it must meet the low latency needs of the network traffic generated by the cache synchronization of Oracle 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.

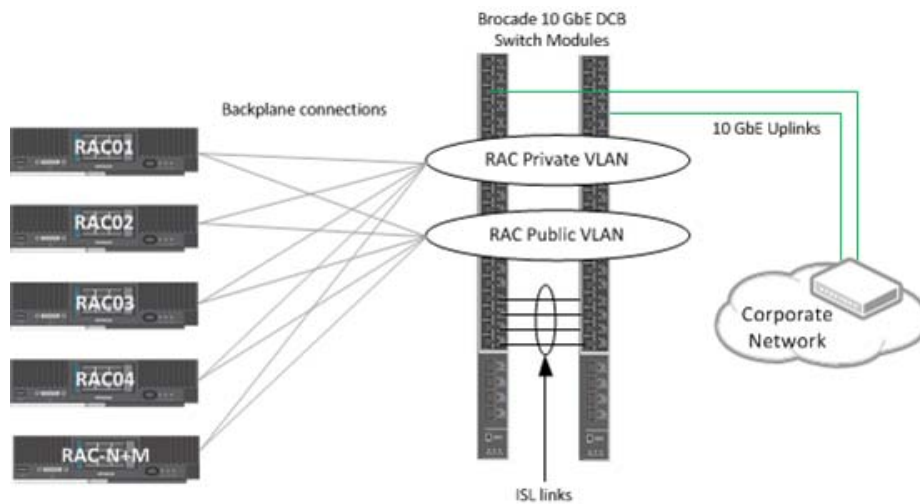


Figure 5

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

Table 11. Network Configuration

Server	NIC Ports	UMC PF Number	Switch Bay ID	Switch Ports (Internal)	VLAN	NIC BOND	Network	Bandwidth (Gb/sec)
RAC01	B1-CNIC-0	0	1	1	1	Bond1	Private	10
	B1-CNIC-1	1	2	1				10
	B1-CNIC-2	2	1	2	2	Bond2	Public Oracle	9
		6	1	2	3	Bond3	Public Management	1
	B1-CNIC-3	3	2	2	2	Bond2	Public Oracle	9
		7	2	2	3	Bond3	Public Management	1

Table 11. Network Configuration (Continued)

Server	NIC Ports	UMC PF Number	Switch Bay ID	Switch Ports (Internal)	VLAN	NIC BOND	Network	Bandwidth (Gb/sec)
RAC02	B3-CNIC-0	0	1	3	1	Bond1	Private	10
	B3-CNIC-1	1	2	3				10
	B3-CNIC-2	2	1	4	2	Bond2	Public Oracle	9
		6	1	4	3	Bond3	Public Management	1
	B3-CNIC-3	3	2	4	2	Bond2	Public Oracle	9
		7	2	4	3	Bond3	Public Management	1
RAC03	B5-CNIC-0	0	1	5	1	Bond1	Private	10
	B5-CNIC-1	1	2	5				10
	B5-CNIC-2	2	1	6	2	Bond2	Public Oracle	9
		6	1	6	3	Bond3	Public Management	1
	B5-CNIC-3	3	2	6	2	Bond2	Public Oracle	9
		7	2	6	3	Bond3	Public Management	1
RAC04	B7-CNIC-0	0	1	7	1	Bond1	Private	10
	B7-CNIC-1	1	2	7				10
	B7-CNIC-2	2	1	8	2	Bond2	Public Oracle	9
		6	1	8	3	Bond3	Public Management	1
	B7-CNIC-3	3	2	8	2	Bond2	Public Oracle	9
		7	2	8	3	Bond3	Public Management	1
RAC-N+M	B9-CNIC-0	0	1	9	1	From failed server	Private	10
	B9-CNIC-1	1	2	9				10
	B9-CNIC-2	2	1	10	2		Public Oracle	9
		6	1	10	3		Public Management	1
	B9-CNIC-3	3	2	10	2		Public Oracle	9
		7	2	10	3		Public Management	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 [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.

 **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

© Hitachi Data Systems Corporation 2016. 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. Microsoft, Windows Server, and Hyper-V are trademarks or registered trademarks of Microsoft Corporation. All other trademarks, service marks, and company names are properties of their respective owners.

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

AS-529-01. September 2016.