Oracle Diamond Partner

HITACHI
Inspire the Next

Hitachi Unified Compute Platform 6000 for Oracle Real Application Clusters on Two Single-CPU Nodes Using Hitachi Virtual Storage Platform F800, Hitachi Accelerated Flash, and Hitachi Compute Blade 2500 With Intel Xeon E5-2699 v4

Reference Architecture Guide

TECHNICAL PAPER

By Eduardo Freitas, Libo Jiang, Steve Burns

January 2017

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 Single-CPU Nodes Using Hitachi Virtual Storage Platform F800, Hitachi Accelerated Flash, and Hitachi Compute Blade 2500 With Intel Xeon E5-2699 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 2-node solution provides a high performance, integrated solution for an Oracle infrastructure. The environment uses Hitachi Virtual Storage Platform F800 (VSP F800) 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 real world workloads such as OLTP and DSS. The dedicated servers run Oracle Database 12c R1 with the Oracle Real Application Cluster (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 F800 (VSP F800) using Hitachi Accelerated Flash (HAF)
- Hitachi Compute Blade 2500 (CB 2500) using 520H B4 server blades with Intel Xeon E5-2699 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
- 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 (UCP for Oracle RAC) on two nodes using Hitachi Virtual Storage Platform F800 (F800) with Hitachi Accelerated Flash (HAF). In this configuration, Hitachi Compute Blade 2500 (CB 2500) uses 520H B4 server blades with Intel Xeon E5-2699 v4 processor CPUs.

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

This reference architecture includes the following:

- **Hitachi Compute Blade 2500** with three server blades
  - **Server Blade 1** — Oracle RAC NODE 1
  - **Server Blade 2** — Oracle RAC NODE 2
  - **Server Blade 3** — N+M Cold Standby
- **Hitachi Virtual Storage Platform F800** with Hitachi Accelerated Flash
- 16 Gb/sec direct-connect SAN infrastructure
- 10 GbE LAN infrastructure

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

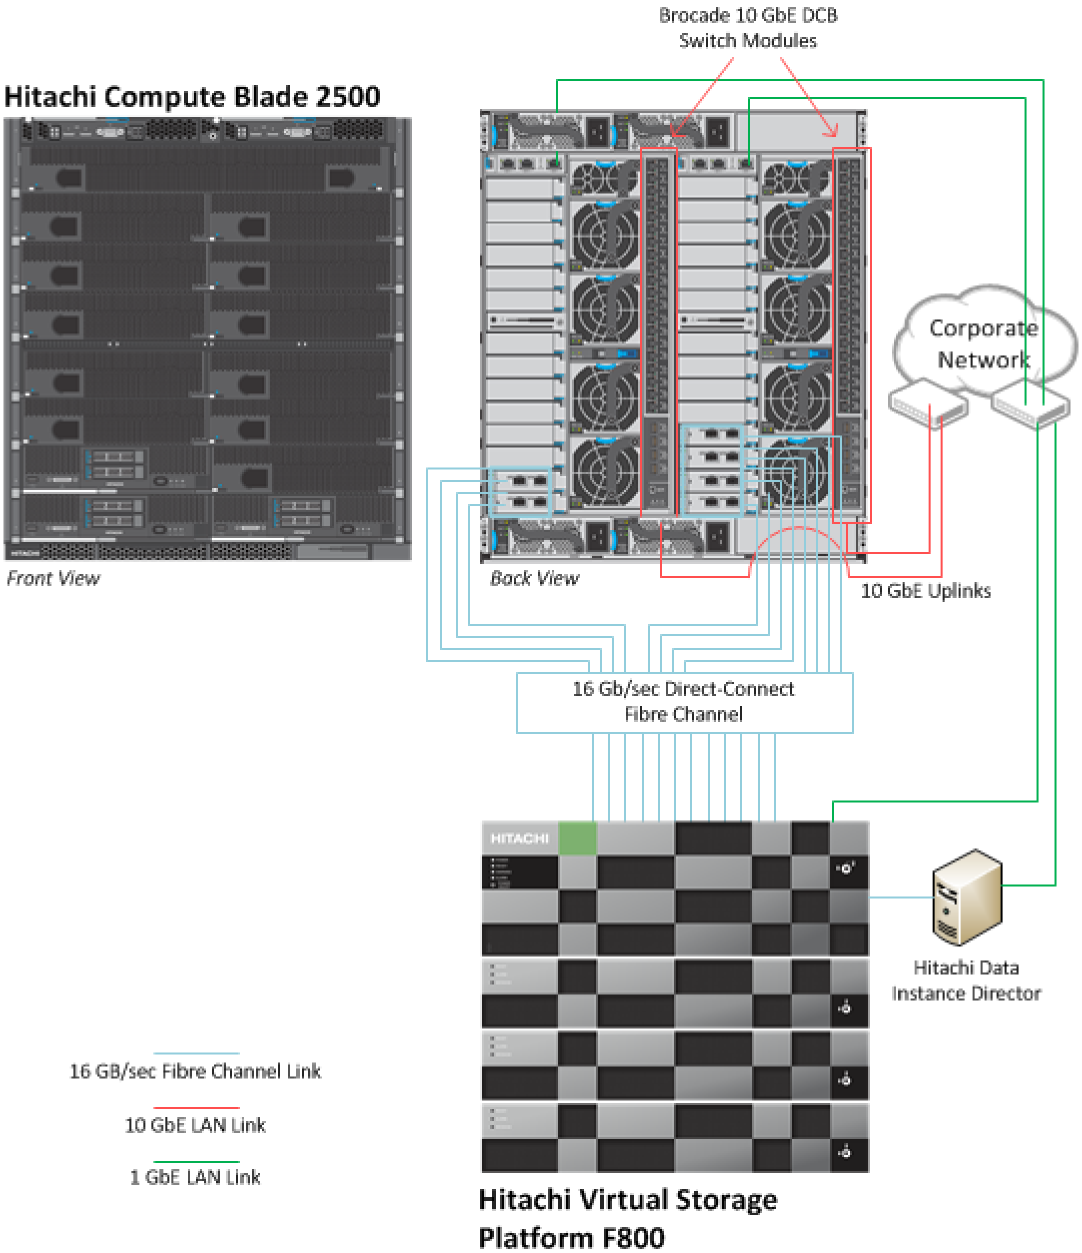

Figure 1

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 F800 (VSP F800)	■ Dual controller ■ 24 × 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
	3.2 TB flash memory drives (FMDs) ■ 2 spares		42 - 11 on the first two trays, and 10 on the last two trays
Hitachi Compute Blade 2500 (CB 2500) chassis	■ 2 × 10 GbE data center bridging (DCB) local area network (LAN) switch module Brocade VDX 2746 ■ 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
520H B4 Half-width Server Blade	■ 1 Intel Xeon E5-2699 v4 processor CPU ■ 192 RAM, 12 × 16 GB DDR4 memory ■ 1 × 4-port 10 GbE converged network adapter (CNA) LAN on motherboard (LOM)	10-04/10-04	3
	■ Hitachi 16 Gb/sec 2-port PCIe Fibre Channel HBA	40-05-00	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 (SVOS) with the following: ▪ Hitachi Storage Navigator (SN) ▪ 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 ASM	12c Release 1 12.1.0.2.0	Oracle ASM
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.

Hitachi Virtual Storage Platform Fx00 Models

Hitachi Virtual Storage Platform Fx00 models deliver superior all-flash performance for business-critical applications, with continuous data availability. High-performance network attached storage with nondisruptive deduplication reduces the required storage capacity by up to 90% with the power to handle large, mixed-workload environments.

Hitachi Storage Virtualization Operating System provides storage virtualization, high availability, superior performance, and advanced data protection for all Virtual Storage Platform Fx00 models. This proven, mature software provides common features to consolidate assets, reclaim space, extend life, and reduce migration effort.

This solution uses Hitachi Virtual Storage Platform F800, 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.

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

Hitachi Storage Virtualization Operating System

Hitachi Storage Virtualization Operating System 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).

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 extended 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.

Hitachi Data Instance Director

Hitachi Data Instance Director (HDID) 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 using the following Brocade products:

- Brocade VDX 2746 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 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 F800 with Hitachi Accelerated Flash.

Specific infrastructure configuration details include the following:

- **Server** — Two server nodes configured in an Oracle RAC with one additional N+M cold standby node.
- **Storage System** — There are LDEVs mapped to each port that are presented to the server as LUNs.
- **SAN Connection** — Each Fibre Channel HBA port is directly connected to the storage front-end ports.

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 diagram of this solution.

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

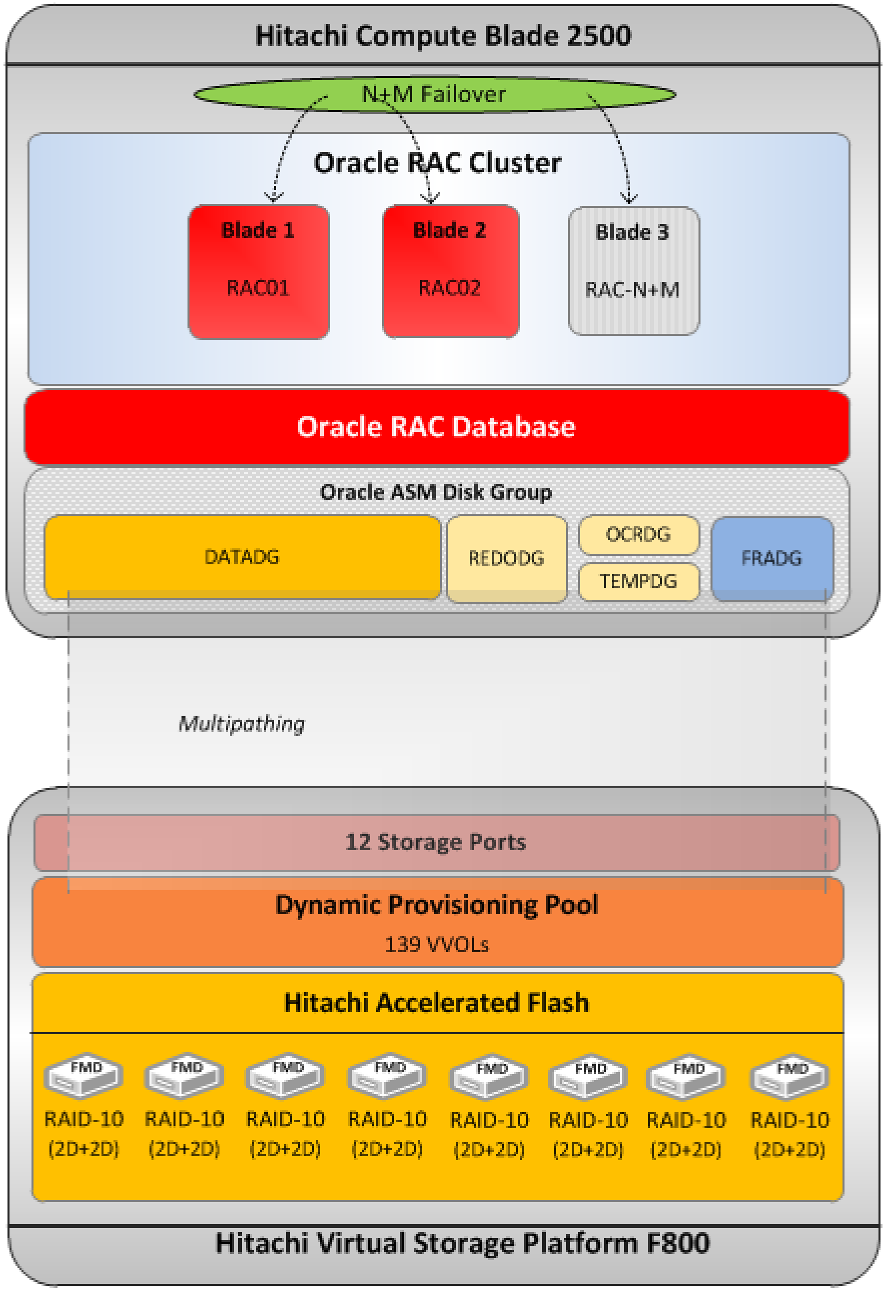

Figure 2

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 (Oracle Production)	ora_bk_fmd_02 (Snapshot Target)
Pool Type	Dynamic Provisioning	Dynamic Provisioning
RAID Group	1-1 – 1-8	1-9
RAID Level	RAID-10 (2D+2D)	RAID-6 (6D+2P)
Number of LDEVs	256 × 204.7 GB	32 × 614.7 GB
Drive Type	3.2 TB Flash Module Drive (FMD)	3.2 TB Flash Module Drive (FMD)
Number of Drives	32, with 8 per tray	8, with 2 per tray
Pool Capacity	51.17 TB	19.19 TB

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

Table 4. Logical Storage Configuration

Pool ID	ora_prod_fmd_01						
Number of VVOLs	3	64	8	32	32	4	1
VVOL Size	5 GB	300 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	12 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 F800 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 a Data ASM diskgroup for the data and index tablespaces. The smallfile tablespace consists of 2048 datafiles 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 small TEMP tablespace. Limit the size of each tempfile to 31 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. 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 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 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 Setting	Use This Value
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 lists 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/sddlmab1	3	▪ Oracle Cluster Registry ▪ Voting Disk
	OCRDISK2	/dev/sddlmac1		
	OCRDISK3	/dev/sddlmad1		
REDODG	RGDISK01	/dev/sddlmae1	32	▪ 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/sddlmbe1		
	RGDISK18	/dev/sddlmbf1		
	RGDISK19	/dev/sddlmbg1		
	RGDISK20	/dev/sddlmbh1		
	RGDISK21	/dev/sddlmbi1		
	RGDISK22	/dev/sddlmbj1		
	RGDISK23	/dev/sddlmbk1		
	RGDISK24	/dev/sddlmbl1		
	RGDISK25	/dev/sddlmbm1		
	RGDISK26	/dev/sddlmbn1		
	RGDISK27	/dev/sddlmbo1		
	RGDISK28	/dev/sddlmbp1		
	RGDISK29	/dev/sddlmca1		
	RGDISK30	/dev/sddlmcb1		
	RGDISK31	/dev/sddlmcc1		

Table 7. Oracle ASM Disk Configuration (Continued)

ASM Disk Group	ASM Disk	HDLM LUNs	LUNs Count	Purpose
	RGDISK32	/dev/sddlmcd1		
DATADG	DADISK01	/dev/sddlmce1	64	▪ Application Data ▪ Undo ▪ System ▪ Sysaux
	DADISK02	/dev/sddlmcf1		
	DADISK03	/dev/sddlmcg1		
	DADISK04	/dev/sddlmch1		
	DADISK05	/dev/sddlmci1		
	DADISK06	/dev/sddlmcj1		
	DADISK07	/dev/sddlmck1		
	DADISK08	/dev/sddlmcl1		
	DADISK09	/dev/sddlmcm1		
	DADISK10	/dev/sddlmcn1		
	DADISK11	/dev/sddlmco1		
	DADISK12	/dev/sddlmcp1		
	DADISK13	/dev/sddlmda1		
	DADISK14	/dev/sddlmdb1		
	DADISK15	/dev/sddlmdc1		
	DADISK16	/dev/sddlmdd1		
	DADISK17	/dev/sddlmde1		
	DADISK18	/dev/sddlmdf1		
	DADISK19	/dev/sddlmdg1		
	DADISK20	/dev/sddlmdh1		
	DADISK21	/dev/sddlmdi1		
	DADISK22	/dev/sddlmdj1		
	DADISK23	/dev/sddlmdk1		
	DADISK24	/dev/sddlmdl1		
	DADISK25	/dev/sddlmdm1		
	DADISK26	/dev/sddlmdn1		
	DADISK27	/dev/sddlmdo1		
	DADISK28	/dev/sddlmdp1		
	DADISK29	/dev/sddlmea1		
	DADISK30	/dev/sddlmeb1		
	DADISK31	/dev/sddlmec1		
	DADISK32	/dev/sddlmed1		
	DADISK33	/dev/sddlmee1		
	DADISK34	/dev/sddlmef1		
	DADISK35	/dev/sddlmeg1		
	DADISK36	/dev/sddlmeh1		

Table 7. Oracle ASM Disk Configuration (Continued)

ASM Disk Group	ASM Disk	HDLM LUNs	LUNs Count	Purpose
	DADISK37	/dev/sddlmei1		
	DADISK38	/dev/sddlmej1		
	DADISK39	/dev/sddlmek1		
	DADISK40	/dev/sddlmel1		
	DADISK41	/dev/sddlmem1		
	DADISK42	/dev/sddlmen1		
	DADISK43	/dev/sddlmeo1		
	DADISK44	/dev/sddlmep1		
	DADISK45	/dev/sddlmfa1		
	DADISK46	/dev/sddlmfb1		
	DADISK47	/dev/sddlmfc1		
	DADISK48	/dev/sddlmfd1		
	DADISK49	/dev/sddlmfe1		
	DADISK50	/dev/sddlmff1		
	DADISK51	/dev/sddlmfg1		
	DADISK52	/dev/sddlmfh1		
	DADISK53	/dev/sddlmfi1		
	DADISK54	/dev/sddlmfj1		
	DADISK55	/dev/sddlmfk1		
	DADISK56	/dev/sddlmfl1		
	DADISK57	/dev/sddlmfm1		
	DADISK58	/dev/sddlmfn1		
	DADISK59	/dev/sddlmfo1		
	DADISK60	/dev/sddlmfp1		
	DADISK61	/dev/sddlmga1		
	DADISK62	/dev/sddlmgb1		
	DADISK63	/dev/sddlmgc1		
	DADISK64	/dev/sddlmgd1		
TEMPDG	TEDISK01	/dev/sddlmge1	8	▪ TEMP
	TEDISK02	/dev/sddlmgf1		
	TEDISK03	/dev/sddlmgg1		
	TEDISK04	/dev/sddlmgh1		
	TEDISK05	/dev/sddlmgi1		
	TEDISK06	/dev/sddlmgj1		
	TEDISK07	/dev/sddlmgk1		
	TEDISK08	/dev/sddlmgl1		

Table 7. Oracle ASM Disk Configuration (Continued)

ASM Disk Group	ASM Disk	HDLM LUNs	LUNs Count	Purpose
FRADG	FRDISK01	/dev/sddlmgm1	32	▪ Archive Logs ▪ Incremental Backups ▪ Control File Autobackups
	FRDISK02	/dev/sddlmgn1		
	FRDISK03	/dev/sddlmgo1		
	FRDISK04	/dev/sddlmgp1		
	FRDISK05	/dev/sddlmha1		
	FRDISK06	/dev/sddlmhb1		
	FRDISK07	/dev/sddlmhc1		
	FRDISK08	/dev/sddlmhd1		
	FRDISK09	/dev/sddlmhe1		
	FRDISK10	/dev/sddlmhf1		
	FRDISK11	/dev/sddlmhg1		
	FRDISK12	/dev/sddlmhh1		
	FRDISK13	/dev/sddlmhi1		
	FRDISK14	/dev/sddlmhj1		
	FRDISK15	/dev/sddlmhk1		
	FRDISK16	/dev/sddlmhl1		
	FRDISK17	/dev/sddlmhm1		
	FRDISK18	/dev/sddlmhn1		
	FRDISK19	/dev/sddlmho1		
	FRDISK20	/dev/sddlmhp1		
	FRDISK21	/dev/sddlmia1		
	FRDISK22	/dev/sddlmib1		
	FRDISK23	/dev/sddlmic1		
	FRDISK24	/dev/sddlmid1		
	FRDISK25	/dev/sddlmie1		
	FRDISK26	/dev/sddlmif1		
	FRDISK27	/dev/sddlmig1		
	FRDISK28	/dev/sddlmih1		
	FRDISK29	/dev/sddlmii1		
	FRDISK30	/dev/sddlmij1		
	FRDISK31	/dev/sddlmik1		
	FRDISK32	/dev/sddlmil1		

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 three server blades.

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

Blade	Form Size	Server Name	Role	CPU Core	RAM
1	Half-Width	RAC01	RAC Node	22	192 GB
2	Half-Width	RAC02	RAC Node	22	192 GB
3	Half-Width	RAC-N+M	Standby	22	192 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 3 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.

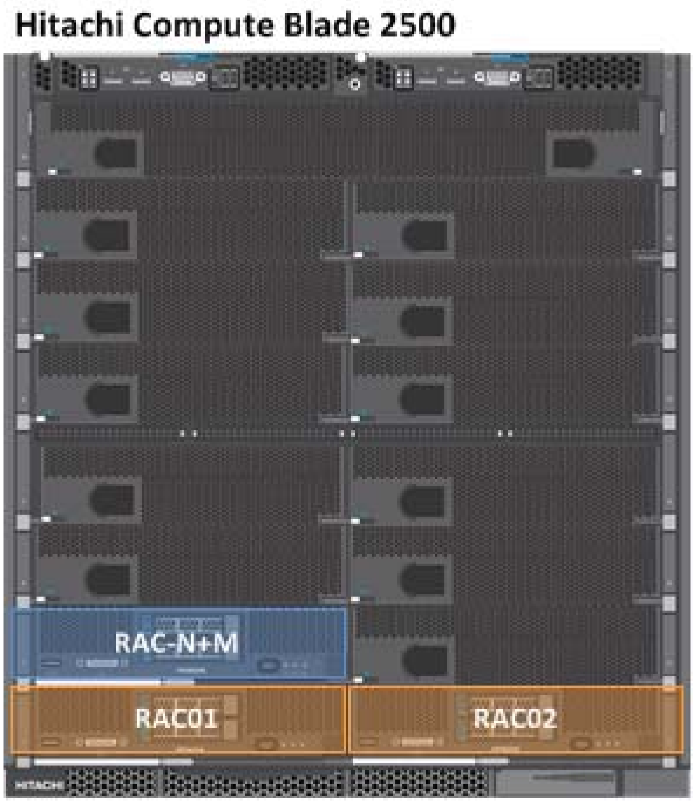

Figure 3

SAN Architecture

Map the provisioned LDEVs to multiple ports on Hitachi Virtual Storage Platform F800 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 5. This provides a four path connection for all LUNs mapped to each of the database servers in Oracle RAC database. Table 9 shows the direct-connect from the HBA of the server blade to the Hitachi Virtual Storage Platform F800 ports.

Table 9. Fibre Channel Direct-Connect Configuration on Hitachi Virtual Storage Platform F800

Host	HBA	Storage Port	Storage Host Group
RAC01	HBA1-1	1A	CB2500_B1_1A
	HBA1-2	2A	CB2500_B1_2A
	HBA2-1	1B	CB2500_B1_1B
	HBA2-2	2B	CB2500_B1_2B
RAC02	HBA1-1	3A	CB2500_B2_3A
	HBA1-2	4A	CB2500_B2_4A
	HBA2-1	3B	CB2500_B2_3B
	HBA2-2	4B	CB2500_B2_4B

Table 9. Fibre Channel Direct-Connect Configuration on Hitachi Virtual Storage Platform F800 (Continued)

Host	HBA	Storage Port	Storage Host Group
RAC-N+M	HBA1-1	1C	CB2500_B1_1A CB2500_B2_3A
	HBA1-2	2C	CB2500_B1_2A CB2500_B2_4A
	HBA2-1	1D	CB2500_B1_1B CB2500_B2_3B
	HBA2-2	2D	CB2500_B1_2B CB2500_B2_4B

Hitachi 16 Gb/sec PCIe HBA Card Configuration

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

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

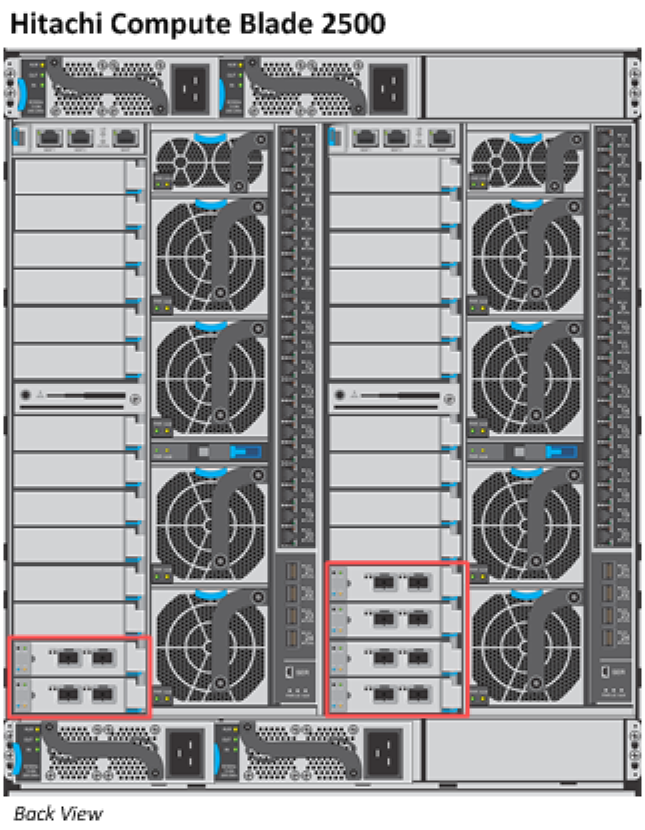

Figure 4

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

Table 10. Hitachi HBA PCIe Card Parameters

For This	Use This
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 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 Gb/sec NICs for the cluster interconnect and public network.

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

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

- For each server in the Oracle RAC clusterware configuration, use at least two identical, high bandwidth, low-latency NICs for the interconnection.
- Use NIC bonding to provide fail over and load balancing of interconnections within a server.
- Set all NICs to full duplex mode.
- Use at least two public NICs for client connections to the application and database.
- Use at least two private NICs for the cluster interconnection.

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

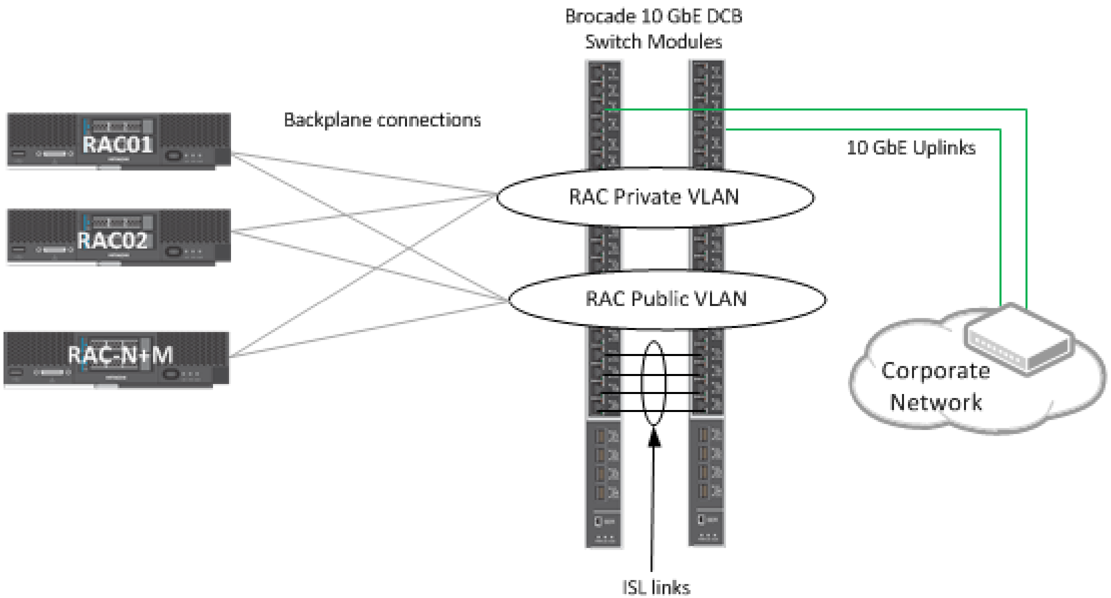

Figure 5

Table 11 shows the network configuration for this solution. Configure the VLAN accordingly 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	15	2	Bond2	Public Oracle	9
		6	1	15	3	Bond3	Public Management	1
	B1-CNIC-3	3	2	15	2	Bond2	Public Oracle	9
		7	2	15	3	Bond3	Public Management	1
RAC02	B2-CNIC-0	0	1	2	1	Bond1	Private	10
	B2-CNIC-1	1	2	2				10
	B2-CNIC-2	2	1	16	2	Bond2	Public Oracle	9
		6	1	16	3	Bond3	Public Management	1
	B2-CNIC-3	3	2	16	2	Bond2	Public Oracle	9
		7	2	16	3	Bond3	Public Management	1

Table 11. Network Configuration

Server	NIC Ports	UMC PF Number	Switch Bay ID	Switch Ports (Internal)	VLAN	NIC BOND	Network	Bandwidth (Gb/sec)
RAC-N+M	B3-CNIC-0	0	1	5	1	From failed server	Private	10
	B3-CNIC-1	1	2	5				10
	B3-CNIC-2	2	1	19	2		Public Oracle	9
		6	1	19	3		Public Management	1
	B3-CNIC-3	3	2	19	2		Public Oracle	9
		7	2	19	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 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-538-02.January 2017.