ORACLE Diamond Partner

HITACHI Inspire the Next

TECHNICAL PAPER

# Hitachi Unified Compute Platform 6000 for Oracle Real Application Clusters on Four Nodes Using Hitachi Virtual Storage Platform F800, Hitachi Accelerated Flash, and Hitachi Compute Blade 2500

## Reference Architecture Guide

By Libo Jiang

September 2016

Hitachi Data Systems

# Feedback

Hitachi Data Systems welcomes your feedback. Please share your thoughts by sending an email message to SolutionLab@hds.com. To assist the routing of this message, use the paper number in the subject and the title of this white paper in the text.

**Note** — This document is HDS Confidential. It is only for use by employees and contracts at a Distribution Center. Proprietary information.

# Contents


Solution Overview ..... 2
Key Solution Components ..... 4
Hitachi Compute Blade 2500 ..... 5
Hitachi Virtual Storage Platform F800 ..... 5
Hitachi Accelerated Flash ..... 5
Hitachi Storage Virtualization Operating System ..... 6
Hitachi Dynamic Link Manager Advanced ..... 6
Brocade Switches ..... 6
Red Hat Enterprise Linux ..... 6
Oracle Database ..... 6
Solution Design ..... 7
Storage Architecture ..... 7
Database Layout ..... 10
Server and Application Architecture ..... 17
SAN Architecture ..... 18
Network Architecture ..... 20
Test Results ..... 23

# Hitachi Unified Compute Platform 6000 for Oracle Real Application Clusters on Four Nodes Using Hitachi Virtual Storage Platform F800, Hitachi Accelerated Flash, and Hitachi Compute Blade 2500

## 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 F800 (VSP F800) with Hitachi Accelerated Flash (HAF). Use this document to design an infrastructure for your requirements and budget.

This validated solution integrates servers, storage systems, network, and storage software. The environment provides reliability, high availability, scalability, and performance while processing small-scale to large-scale OLTP workloads. The dedicated servers run 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 F800 (VSP F800) using Hitachi Accelerated Flash (HAF)
- Hitachi Compute Blade 2500 (CB 2500)
- Storage area networks
- Oracle RAC Database 12c Release 1
- Oracle Automatic Storage Management (ASM)
- 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 four nodes using Hitachi Virtual Storage Platform F800 (VSP F800) with Hitachi Accelerated Flash (HAF). This environment addresses the high availability, performance, and scalability requirements for OLTP workloads. Tailor your implementation of this solution to meet your specific needs.

This reference architecture includes the following:

- **Hitachi Compute Blade 2500** with four server blades
  - **Server Blade 1** — Oracle RAC NODE 1
  - **Server Blade 2** — Oracle RAC NODE 2
  - **Server Blade 3** — Oracle RAC NODE 3
  - **Server Blade 4** — Oracle RAC NODE 4
- **Hitachi Virtual Storage Platform F800** with Hitachi Accelerated Flash
- 16 Gb/sec SAN infrastructure
- 10 GbE LAN infrastructure

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

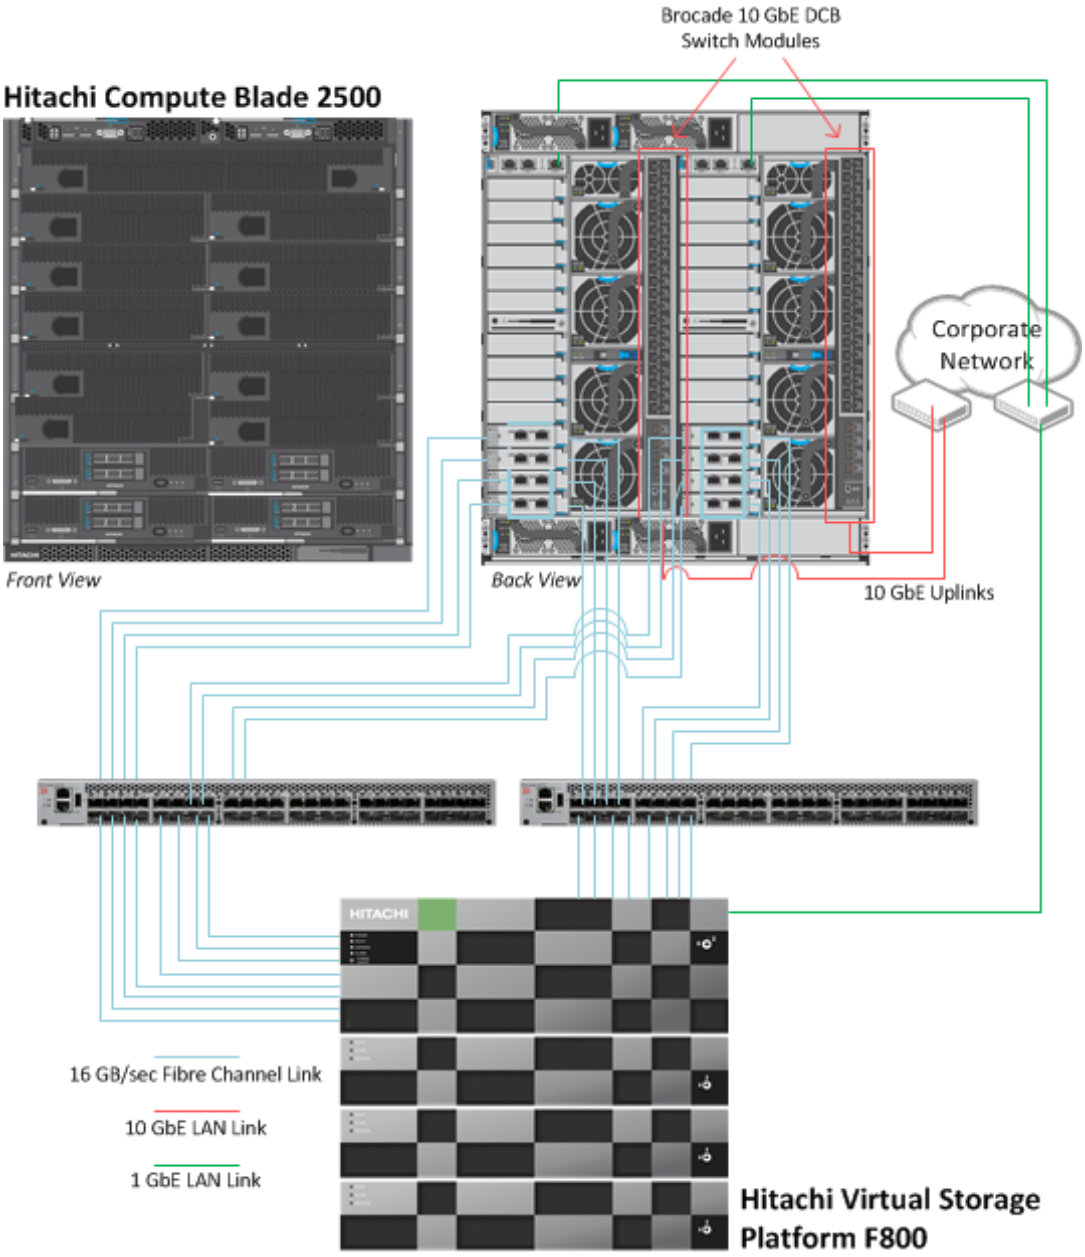


**Figure 1**

# 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<br>16 × 16 Gb/sec Fibre Channel ports<br>8 × 12 Gb/sec backend serial attached SCSI (SAS) ports<br>512 GB cache memory | 83-03-23-60/00 | 1 |
| | 1.6 TB flash memory drives (FMDs)<br>■ 2 spares | | 42 |
| | | | |
| Hitachi Compute Blade 2500 (CB 2500) chassis | 2 × 10 GbE data center bridging (DCB) local area network (LAN) switch module<br>10 Fan modules<br>2 Management modules<br>4 Power supply modules | Management Module Firmware Version A0150-B-1410<br>Dictionary Version A0027<br>DCB Switch Version 4.0.1_hit1 | 1 |
| 520H B3 Half-width Server Blade | 2 Intel Xeon E5-2699 v3 CPU<br>256 RAM, 16 × 16 GB DDR4 memory<br>1 × 4-port 10 GbE converged network adapter (CNA) LAN on motherboard (LOM) | 08-46 | 4 |
| | Hitachi 16 Gb/sec 2-port PCIe Fibre Channel HBA | 40-04-00 | 8 |
| Brocade 6510 Fibre Channel switch | ■ 48 port Fibre Channel switch<br>■ 16 × 16 Gb/sec SFPs | 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 Navigator | N/A | Storage management suite |
| Hitachi Command Suite | 8 | Storage management suite |
| | | |
| 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 |
| Hitachi Dynamic Link Manager Advanced | 8.4.0-00 | Multipath software |
| Hitachi Dynamic Provisioning | Virtual Storage Platform F800 | Storage license |
| Flash Optimization Feature | Virtual Storage Platform F800 | Storage license |
| Hitachi Accelerated Flash | Virtual Storage Platform F800 | Storage license |

## 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 F800

Hitachi Virtual Storage Platform F800 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 F800 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 F800 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 Virtual Storage Platform F800 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 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 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 VDX 2746, a 10 GbE switch module
- Brocade 6510 48 port Fibre Channel switch

## Red Hat Enterprise Linux

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

## Oracle Database

Oracle Database has a multitenant architecture so you can consolidate many databases quickly and manage them as a cloud service. Oracle Database also includes in-memory data processing capabilities for analytical performance. Additional database innovations deliver efficiency, performance, security, and availability. Oracle Database comes in two editions: Enterprise Edition and Standard Edition 2.

Oracle Real Application Clusters (Oracle RAC) is a clustered version of Oracle Database. It is based on a comprehensive high-availability stack that can be used as the foundation of a database cloud system, as well as a shared infrastructure. This ensures high availability, scalability, and agility for any application.

Oracle Automatic Storage Management (Oracle ASM) is a volume manager and a file system for Oracle database files. This supports single-instance Oracle Database and Oracle RAC configurations. Oracle ASM is the recommended storage management solution that provides an alternative to conventional volume managers, file systems, and raw devices.

# Solution Design

This describes the reference architecture environment, implementing a half-rack environment for Hitachi Unified Compute Platform 6000 for Oracle Real Application Cluster using Hitachi Virtual Storage Platform F800 with Hitachi Accelerated Flash.

Specific infrastructure configuration details include the following:

- **Server** — Four server nodes configured in an Oracle Real Application Cluster.
- **Storage System** — There are LDEVs mapped to each port that are presented to the server as LUNs.
- **SAN Connection** — Connect each Fibre Channel HBA port to the storage front-end port 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 diagram of this solution.

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

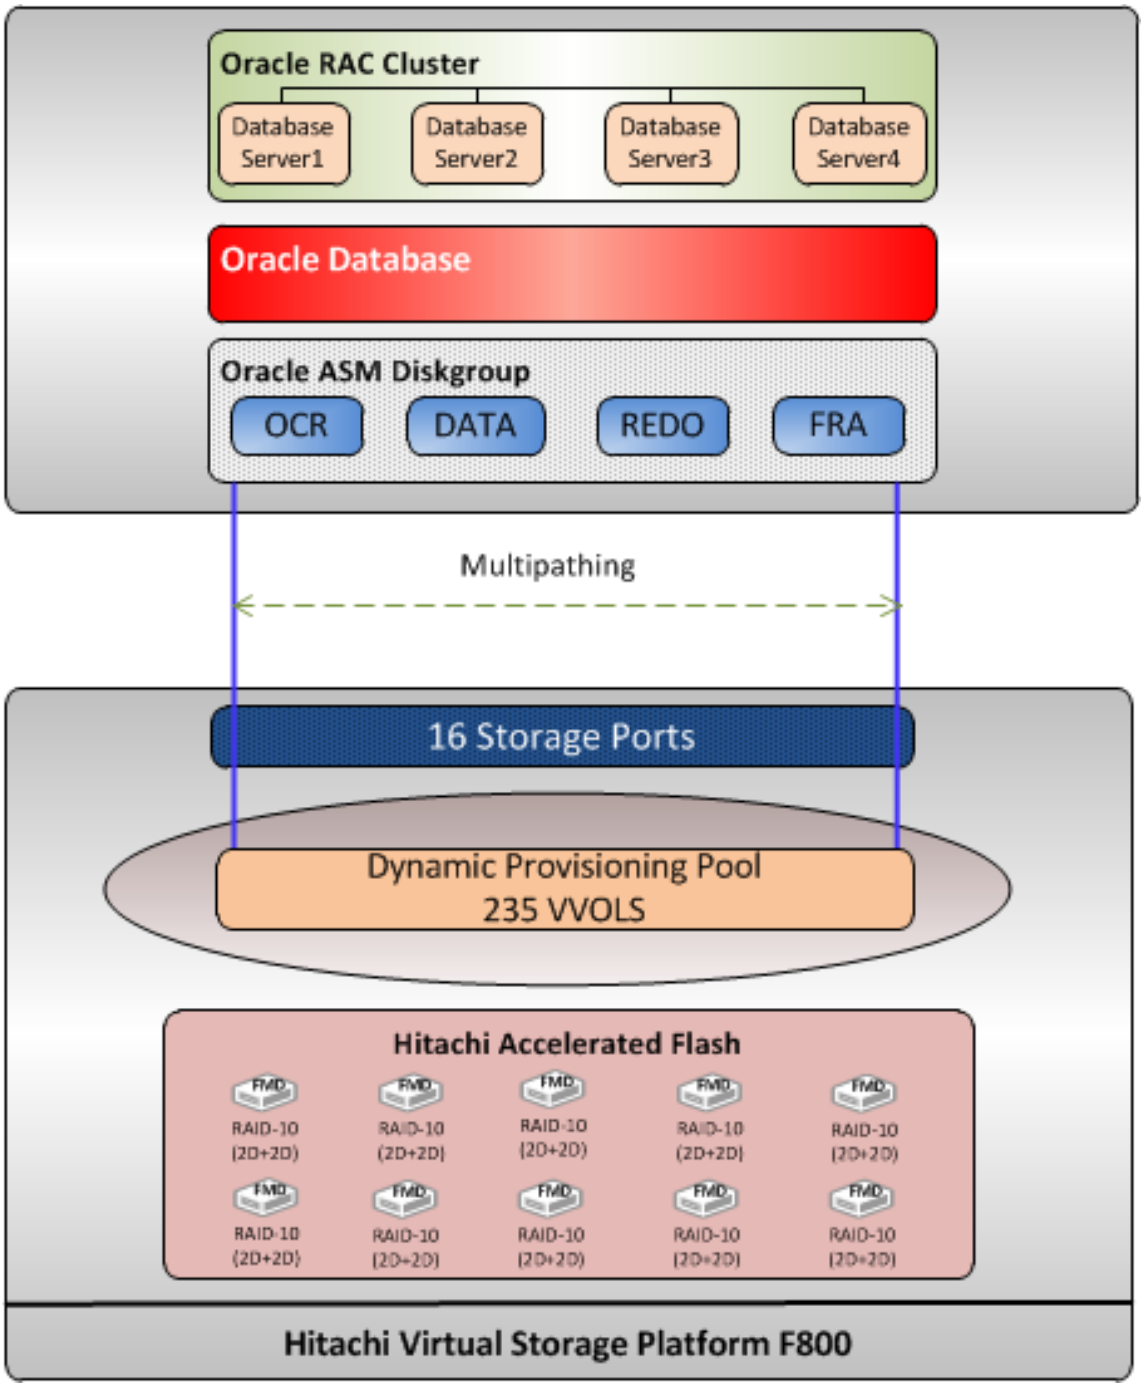


**Figure 2**

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

**Table 3. Storage Pool Configuration**

| | |
|---|---|
| **Pool ID** | ora_dp_fmd_01 |
| **Pool Type** | Dynamic Provisioning |
| **RAID Group** | 1-1 – 1-10 |
| **RAID Level** | RAID-10 (2D+2D) |
| **Number of LDEVs** | 320 × 102.3 GB |
| **Drive Type** | 1.6 TB Flash Module Drive (FMD) |
| **Number of Drives** | 16, with 4 per tray |
| **Pool Capacity** | 31.96 TB |

Table 4 shows the logical storage configuration used in the tested configuration. You may use a different configuration.

**Table 4. Logical Storage Configuration**

| | | | | | |
|---|---|---|---|---|---|
| **Pool ID** | ora_dp_fmd_01 | | | | |
| **Number of VVOLs** | 160 | 32 | 40 | 3 | 4 |
| **VVOL Size** | 120 GB | 10 GB | 300 GB | 5 GB | 200 GB |
| **ASM Disk group** | DATA | REDO | FRA | OCR | |
| **Purpose** | OLTP Application Tablespaces<br>System<br>Sysaux<br>Undo<br>Temp | Online Redo Logs<br>Control Files | Incremental Backups<br>Archived Redo Logs<br>Control File Autobackups | Oracle Cluster Registry<br>Voting Disk | SAN Boot |
| **Storage Port** | 1A, 1B, 1C, 1D, 2A, 2B, 2C, 2D, 3A, 3B, 3C, 3D, 4A, 4B,4C,4D | | | | |

## 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 disk group for the data and index tablespaces. The smallfile tablespace for data consists of 2048 datafiles that are 8 GB each.
- **TEMP Tablespace** — Create a bigfile temporary tablespace from a Data ASM disk group in this configuration.
- **Undo Tablespace** — Create four bigfile UNDO tablespaces from the Data ASM disk group. Assign one UNDO tablespace for each database instance in this 4-node Oracle RAC database.
- **Online Redo Logs** — Assign an ASM disk group REDO for online redo logs. 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 the OCR ASM disk group in this 4-node Oracle RAC configuration.
- **Size Settings** — Set the database block size to 8 KB. Set the ASM allocation unit to 1 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 | 512 MB |
| 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 |
|---|---|---|---|---|
| OCR | OCRDISK1 | /dev/sddlmab1 | 3 | ■ Oracle Cluster Registry<br>■ Voting Disk |
| | OCRDISK2 | /dev/sddlmac1 | | |
| | OCRDISK3 | /dev/sddlmad1 | | |
| REDO | REDDISK01 | /dev/sddlmae1 | 32 | ■ Online REDO log groups<br>■ Control files |
| | REDDISK02 | /dev/sddlmaf1 | | |
| | REDDISK03 | /dev/sddlmag1 | | |
| | REDDISK04 | /dev/sddlmah1 | | |
| | REDDISK05 | /dev/sddlmai1 | | |
| | REDDISK06 | /dev/sddlmaj1 | | |
| | REDDISK07 | /dev/sddlmak1 | | |
| | REDDISK08 | /dev/sddlmal1 | | |
| | REDDISK09 | /dev/sddlmam1 | | |
| | REDDISK10 | /dev/sddlman1 | | |
| | REDDISK11 | /dev/sddlmao1 | | |
| | REDDISK12 | /dev/sddlmap1 | | |
| | REDDISK13 | /dev/sddlmba1 | | |
| | REDDISK14 | /dev/sddlmbb1 | | |
| | REDDISK15 | /dev/sddlmbc1 | | |
| | REDDISK16 | /dev/sddlmbd1 | | |
| | REDDISK17 | /dev/sddlmbe1 | | |
| | REDDISK18 | /dev/sddlmbf1 | | |
| | REDDISK19 | /dev/sddlmbg1 | | |
| | REDDISK20 | /dev/sddlmbh1 | | |
| | REDDISK21 | /dev/sddlmbi1 | | |
| | REDDISK22 | /dev/sddlmbj1 | | |
| | REDDISK23 | /dev/sddlmbk1 | | |
| | REDDISK24 | /dev/sddlmbl1 | | |
| | REDDISK25 | /dev/sddlmbm1 | | |
| | REDDISK26 | /dev/sddlmbn1 | | |
| | REDDISK27 | /dev/sddlmbo1 | | |
| | REDDISK28 | /dev/sddlmbp1 | | |
| | REDDISK29 | /dev/sddlmca1 | | |
| | REDDISK30 | /dev/sddlmcb1 | | |
| | REDDISK31 | /dev/sddlmcc1 | | |
| | REDDISK32 | /dev/sddlmcd1 | | |

Table 7. Oracle ASM Disk Configuration (Continued)

| ASM Disk Group | ASM Disk | HDLM LUNs | LUNs Count | Purpose |
|---|---|---|---|---|
| DATA | DATDISK01 | /dev/sddlmce1 | 160 | ■ Application Data<br>■ System<br>■ Sysaux<br>■ Undo<br>■ Temp |
| | DATDISK02 | /dev/sddlmcf1 | | |
| | DATDISK03 | /dev/sddlmcg1 | | |
| | DATDISK04 | /dev/sddlmch1 | | |
| | DATDISK05 | /dev/sddlmci1 | | |
| | DATDISK06 | /dev/sddlmcj1 | | |
| | DATDISK07 | /dev/sddlmck1 | | |
| | DATDISK08 | /dev/sddlmcl1 | | |
| | DATDISK09 | /dev/sddlmcm1 | | |
| | DATDISK10 | /dev/sddlmcn1 | | |
| | DATDISK11 | /dev/sddlmco1 | | |
| | DATDISK12 | /dev/sddlmcp1 | | |
| | DATDISK13 | /dev/sddlmda1 | | |
| | DATDISK14 | /dev/sddlmdb1 | | |
| | DATDISK15 | /dev/sddlmdc1 | | |
| | DATDISK16 | /dev/sddlmdd1 | | |
| | DATDISK17 | /dev/sddlmde1 | | |
| | DATDISK18 | /dev/sddlmdf1 | | |
| | DATDISK19 | /dev/sddlmdg1 | | |
| | DATDISK20 | /dev/sddlmdh1 | | |
| | DATDISK21 | /dev/sddlmdi1 | | |
| | DATDISK22 | /dev/sddlmdj1 | | |
| | DATDISK23 | /dev/sddlmdk1 | | |
| | DATDISK24 | /dev/sddlmdl1 | | |
| | DATDISK25 | /dev/sddlmdm1 | | |
| | DATDISK26 | /dev/sddlmdn1 | | |
| | DATDISK27 | /dev/sddlmdo1 | | |
| | DATDISK28 | /dev/sddlmdp1 | | |
| | DATDISK29 | /dev/sddlmea1 | | |
| | DATDISK30 | /dev/sddlmeb1 | | |
| | DATDISK31 | /dev/sddlmec1 | | |
| | DATDISK32 | /dev/sddlmed1 | | |
| | DATDISK33 | /dev/sddlmee1 | | |
| | DATDISK34 | /dev/sddlmef1 | | |
| | DATDISK35 | /dev/sddlmeg1 | | |
| | DATDISK36 | /dev/sddlmeh1 | | |
| | DATDISK37 | /dev/sddlmei1 | | |
| | DATDISK38 | /dev/sddlmej1 | | |

Table 7. Oracle ASM Disk Configuration (Continued)

| ASM Disk Group | ASM Disk | HDLM LUNs | LUNs Count | Purpose |
|---|---|---|---|---|
| DATA (Continued) | DATDISK39 | /dev/sddlmek1 | | |
| | DATDISK40 | /dev/sddlmel1 | | |
| | DATDISK41 | /dev/sddlmem1 | | |
| | DATDISK42 | /dev/sddlmen1 | | |
| | DATDISK43 | /dev/sddlmeo1 | | |
| | DATDISK44 | /dev/sddlmep1 | | |
| | DATDISK45 | /dev/sddlmfa1 | | |
| | DATDISK46 | /dev/sddlmfb1 | | |
| | DATDISK47 | /dev/sddlmfc1 | | |
| | DATDISK48 | /dev/sddlmfd1 | | |
| | DATDISK49 | /dev/sddlmfe1 | | |
| | DATDISK50 | /dev/sddlmff1 | | |
| | DATDISK51 | /dev/sddlmfg1 | | |
| | DATDISK52 | /dev/sddlmfh1 | | |
| | DATDISK53 | /dev/sddlmfi1 | | |
| | DATDISK54 | /dev/sddlmfj1 | | |
| | DATDISK55 | /dev/sddlmfk1 | | |
| | DATDISK56 | /dev/sddlmfl1 | | |
| | DATDISK57 | /dev/sddlmfm1 | | |
| | DATDISK58 | /dev/sddlmfn1 | | |
| | DATDISK59 | /dev/sddlmfo1 | | |
| | DATDISK60 | /dev/sddlmfp1 | | |
| | DATDISK61 | /dev/sddlmga1 | | |
| | DATDISK62 | /dev/sddlmgb1 | | |
| | DATDISK63 | /dev/sddlmgc1 | | |
| | DATDISK64 | /dev/sddlmgd1 | | |
| | DATDISK65 | /dev/sddlmge1 | | |
| | DATDISK66 | /dev/sddlmgf1 | | |
| | DATDISK67 | /dev/sddlmgg1 | | |
| | DATDISK68 | /dev/sddlmgh1 | | |
| | DATDISK69 | /dev/sddlmgi1 | | |
| | DATDISK70 | /dev/sddlmgj1 | | |
| | DATDISK71 | /dev/sddlmgk1 | | |
| | DATDISK72 | /dev/sddlmgl1 | | |
| | DATDISK73 | /dev/sddlmgm1 | | |
| | DATDISK74 | /dev/sddlmgn1 | | |
| | DATDISK75 | /dev/sddlmgo1 | | |
| | DATDISK76 | /dev/sddlmgp1 | | |

Table 7. Oracle ASM Disk Configuration (Continued)

| ASM Disk Group | ASM Disk | HDLM LUNs | LUNs Count | Purpose |
|---|---|---|---|---|
| DATA (Continued) | DATDISK77 | /dev/sddlmha1 | | |
| | DATDISK78 | /dev/sddlmhb1 | | |
| | DATDISK79 | /dev/sddlmhc1 | | |
| | DATDISK80 | /dev/sddlmhd1 | | |
| | DATDISK81 | /dev/sddlmhe1 | | |
| | DATDISK82 | /dev/sddlmhf1 | | |
| | DATDISK83 | /dev/sddlmhg1 | | |
| | DATDISK84 | /dev/sddlmhh1 | | |
| | DATDISK85 | /dev/sddlmhi1 | | |
| | DATDISK86 | /dev/sddlmhj1 | | |
| | DATDISK87 | /dev/sddlmhk1 | | |
| | DATDISK88 | /dev/sddlmhl1 | | |
| | DATDISK89 | /dev/sddlmhm1 | | |
| | DATDISK90 | /dev/sddlmhn1 | | |
| | DATDISK91 | /dev/sddlmho1 | | |
| | DATDISK92 | /dev/sddlmhp1 | | |
| | DATDISK93 | /dev/sddlmia1 | | |
| | DATDISK94 | /dev/sddlmib1 | | |
| | DATDISK95 | /dev/sddlmic1 | | |
| | DATDISK96 | /dev/sddlmid1 | | |
| | DATDISK97 | /dev/sddlmie1 | | |
| | DATDISK98 | /dev/sddlmif1 | | |
| | DATDISK99 | /dev/sddlmig1 | | |
| | DATDISK100 | /dev/sddlmih1 | | |
| | DATDISK101 | /dev/sddlmii1 | | |
| | DATDISK102 | /dev/sddlmij1 | | |
| | DATDISK103 | /dev/sddlmik1 | | |
| | DATDISK104 | /dev/sddlmil1 | | |
| | DATDISK105 | /dev/sddlmim1 | | |
| | DATDISK106 | /dev/sddlmin1 | | |
| | DATDISK107 | /dev/sddlmio1 | | |
| | DATDISK108 | /dev/sddlmip1 | | |
| | DATDISK109 | /dev/sddlmja1 | | |
| | DATDISK110 | /dev/sddlmjb1 | | |
| | DATDISK111 | /dev/sddlmjc1 | | |
| | DATDISK112 | /dev/sddlmjd1 | | |
| | DATDISK113 | /dev/sddlmje1 | | |
| | DATDISK114 | /dev/sddlmjf1 | | |

Table 7. Oracle ASM Disk Configuration (Continued)

| ASM Disk Group | ASM Disk | HDLM LUNs | LUNs Count | Purpose |
|---|---|---|---|---|
| DATA (Continued) | DATDISK115 | /dev/sddlmjg1 | | |
| | DATDISK116 | /dev/sddlmjh1 | | |
| | DATDISK117 | /dev/sddlmji1 | | |
| | DATDISK118 | /dev/sddlmjj1 | | |
| | DATDISK119 | /dev/sddlmjk1 | | |
| | DATDISK120 | /dev/sddlmjl1 | | |
| | DATDISK121 | /dev/sddlmjm1 | | |
| | DATDISK122 | /dev/sddlmjn1 | | |
| | DATDISK123 | /dev/sddlmjo1 | | |
| | DATDISK124 | /dev/sddlmjp1 | | |
| | DATDISK125 | /dev/sddlmma1 | | |
| | DATDISK126 | /dev/sddlmmb1 | | |
| | DATDISK127 | /dev/sddlmmc1 | | |
| | DATDISK128 | /dev/sddlmmd1 | | |
| | DATDISK129 | /dev/sddlmme1 | | |
| | DATDISK130 | /dev/sddlmmf1 | | |
| | DATDISK131 | /dev/sddlmmg1 | | |
| | DATDISK132 | /dev/sddlmmh1 | | |
| | DATDISK133 | /dev/sddlmmi1 | | |
| | DATDISK134 | /dev/sddlmmj1 | | |
| | DATDISK135 | /dev/sddlmmk1 | | |
| | DATDISK136 | /dev/sddlmml1 | | |
| | DATDISK137 | /dev/sddlmmm1 | | |
| | DATDISK138 | /dev/sddlmmn1 | | |
| | DATDISK139 | /dev/sddlmmo1 | | |
| | DATDISK140 | /dev/sddlmmp1 | | |
| | DATDISK141 | /dev/sddlmna1 | | |
| | DATDISK142 | /dev/sddlmnb1 | | |
| | DATDISK143 | /dev/sddlmnc1 | | |
| | DATDISK144 | /dev/sddlmnd1 | | |
| | DATDISK145 | /dev/sddlmne1 | | |
| | DATDISK146 | /dev/sddlmnf1 | | |
| | DATDISK147 | /dev/sddlmng1 | | |
| | DATDISK148 | /dev/sddlmnh1 | | |
| | DATDISK149 | /dev/sddlmni1 | | |
| | DATDISK150 | /dev/sddlmnj1 | | |
| | DATDISK151 | /dev/sddlmnk1 | | |
| | DATDISK152 | /dev/sddlmnl1 | | |

**Table 7. Oracle ASM Disk Configuration (Continued)**

| ASM Disk Group | ASM Disk | HDLM LUNs | LUNs Count | Purpose |
|---|---|---|---|---|
| DATA (Continued) | DATDISK153 | /dev/sddlmnm1 | | |
| | DATDISK154 | /dev/sddlmnn1 | | |
| | DATDISK155 | /dev/sddlmno1 | | |
| | DATDISK156 | /dev/sddlmnp1 | | |
| | DATDISK157 | /dev/sddlmpa1 | | |
| | DATDISK158 | /dev/sddlmpb1 | | |
| | DATDISK159 | /dev/sddlmpc1 | | |
| | DATDISK160 | /dev/sddlmpd1 | | |
| FRA | FRADISK01 | /dev/sddlmpe1 | 40 | ■ Archive Logs<br>■ Incremental Backups<br>■ Control File Autobackups |
| | FRADISK02 | /dev/sddlmpf1 | | |
| | FRADISK03 | /dev/sddlmpg1 | | |
| | FRADISK04 | /dev/sddlmph1 | | |
| | FRADISK05 | /dev/sddlmpi1 | | |
| | FRADISK06 | /dev/sddlmpj1 | | |
| | FRADISK07 | /dev/sddlmpk1 | | |
| | FRADISK08 | /dev/sddlmpl1 | | |
| | FRADISK09 | /dev/sddlmpm1 | | |
| | FRADISK10 | /dev/sddlmpn1 | | |
| | FRADISK11 | /dev/sddlmpo1 | | |
| | FRADISK12 | /dev/sddlmpp1 | | |
| | FRADISK13 | /dev/sddlmaaa1 | | |
| | FRADISK14 | /dev/sddlmaab1 | | |
| | FRADISK15 | /dev/sddlmaac1 | | |
| | FRADISK16 | /dev/sddlmaad1 | | |
| | FRADISK17 | /dev/sddlmaae1 | | |
| | FRADISK18 | /dev/sddlmaaf1 | | |
| | FRADISK19 | /dev/sddlmaag1 | | |
| | FRADISK20 | /dev/sddlmaah1 | | |
| | FRADISK21 | /dev/sddlmaai1 | | |
| | FRADISK22 | /dev/sddlmaaj1 | | |
| | FRADISK23 | /dev/sddlmaak1 | | |
| | FRADISK24 | /dev/sddlmaal1 | | |
| | FRADISK25 | /dev/sddlmaam1 | | |
| | FRADISK26 | /dev/sddlmaan1 | | |
| | FRADISK27 | /dev/sddlmaao1 | | |
| | FRADISK28 | /dev/sddlmaap1 | | |
| | FRADISK29 | /dev/sddlmaba1 | | |
| | FRADISK30 | /dev/sddlmabb1 | | |

**Table 7. Oracle ASM Disk Configuration (Continued)**

| ASM Disk Group | ASM Disk | HDLM LUNs | LUNs Count | Purpose |
|---|---|---|---|---|
| FRA (Continued) | FRADISK31 | /dev/sddlmabc1 | | |
| | FRADISK32 | /dev/sddlmabd1 | | |
| | FRADISK33 | /dev/sddlmabe1 | | |
| | FRADISK34 | /dev/sddlmabf1 | | |
| | FRADISK35 | /dev/sddlmabg1 | | |
| | FRADISK36 | /dev/sddlmabh1 | | |
| | FRADISK37 | /dev/sddlmabi1 | | |
| | FRADISK38 | /dev/sddlmabj1 | | |
| | FRADISK39 | /dev/sddlmabk1 | | |
| | FRADISK40 | /dev/sddlmabl1 | | |

## Server and Application Architecture

This reference architecture uses a single Hitachi Compute Blade 2500 chassis with four server blades.

This provides the compute power for Oracle RAC database to handle complex database queries and a large volume of transaction processing in parallel. Table 8 describes the details of the server configuration for this solution.

**Table 8. Server Details**

| Server | Form Size | Server Name | Role | CPU Core | RAM | Blade |
|---|---|---|---|---|---|---|
| Node1 | Half-Width | Jaguar1 | RAC Node | 36 | 256 GB | 1 |
| Node2 | Half-Width | Jaguar2 | RAC Node | 36 | 256 GB | 2 |
| Node3 | Half-Width | Jaguar3 | RAC Node | 36 | 256 GB | 3 |
| Node4 | Half-Width | Jaguar4 | RAC Node | 36 | 256 GB | 4 |

Figure 3 shows the server infrastructure for the reference architecture.

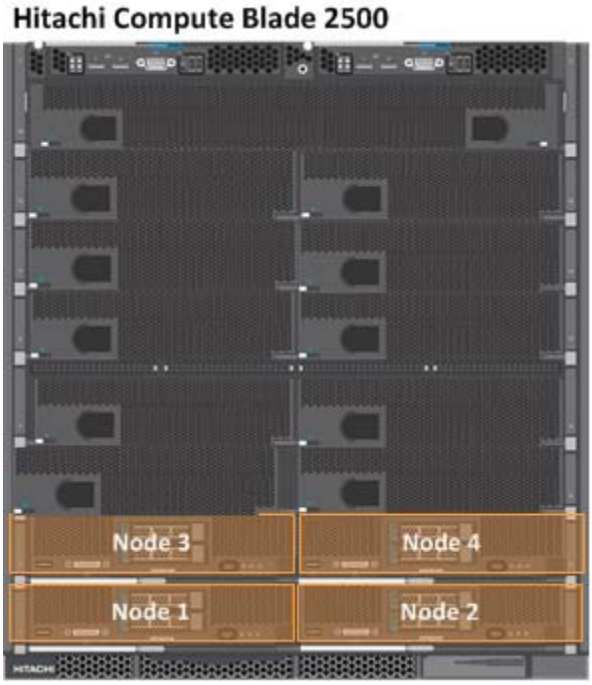


**Figure 3**

## SAN Architecture

Map the provisioned LDEVs to multiple ports on Hitachi Virtual Storage Platform 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, “Key Solution Components From Hitachi Data Systems,” on page 4. This provides a four-path connection for all LUNs mapped to each of the database servers in the Oracle RAC environment. Table 9 shows the SAN connection from the HBA of the server blade to the Hitachi Virtual Storage Platform F800 ports.

**Table 9. Fibre Channel SAN Connect Configuration on Hitachi Virtual Storage Platform F800**

| Host | HBA | Storage Port | Storage Host Group | Zone Name |
|---|---|---|---|---|
| BLADE 1 | HBA1-1 | 1A | BL2500_B1_1A | BL1_HBA1_1_F800_1A |
| | HBA1-2 | 2A | BL2500_B1_2A | BL1_HBA1_2_F800_2A |
| | HBA2-1 | 1B | BL2500_B1_1B | BL1_HBA2_1_F800_1B |
| | HBA2-2 | 2B | BL2500_B1_2B | BL1_HBA2_2_F800_2B |
| BLADE 2 | HBA1-1 | 3A | BL2500_B2_3A | BL2_HBA1_1_F800_3A |
| | HBA1-2 | 4A | BL2500_B2_4A | BL2_HBA1_2_F800_4A |
| | HBA2-1 | 3B | BL2500_B2_3B | BL2_HBA2_1_F800_3B |
| | HBA2-2 | 4B | BL2500_B2_4B | BL2_HBA2_2_F800_4B |

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

| Host | HBA | Storage Port | Storage Host Group | Zone Name |
|---|---|---|---|---|
| BLADE 3 | HBA1-1 | 1C | BL2500_B3_1C | BL3_HBA1_1_F800_1C |
| | HBA1-2 | 2C | BL2500_B3_2C | BL3_HBA1_2_F800_2C |
| | HBA2-1 | 1D | BL2500_B3_1D | BL3_HBA2_1_F800_1D |
| | HBA2-2 | 2D | BL2500_B3_2D | BL3_HBA2_2_F800_2D |
| BLADE 4 | HBA1-1 | 3C | BL2500_B4_3C | BL4_HBA1_1_F800_3C |
| | HBA1-2 | 4C | BL2500_B4_4C | BL4_HBA1_2_F800_4C |
| | HBA2-1 | 3D | BL2500_B4_3D | BL4_HBA2_1_F800_3D |
| | HBA2-2 | 4D | BL2500_B4_4D | BL4_HBA2_2_F800_4D |

## 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 Hitachi Compute Blade 2500.

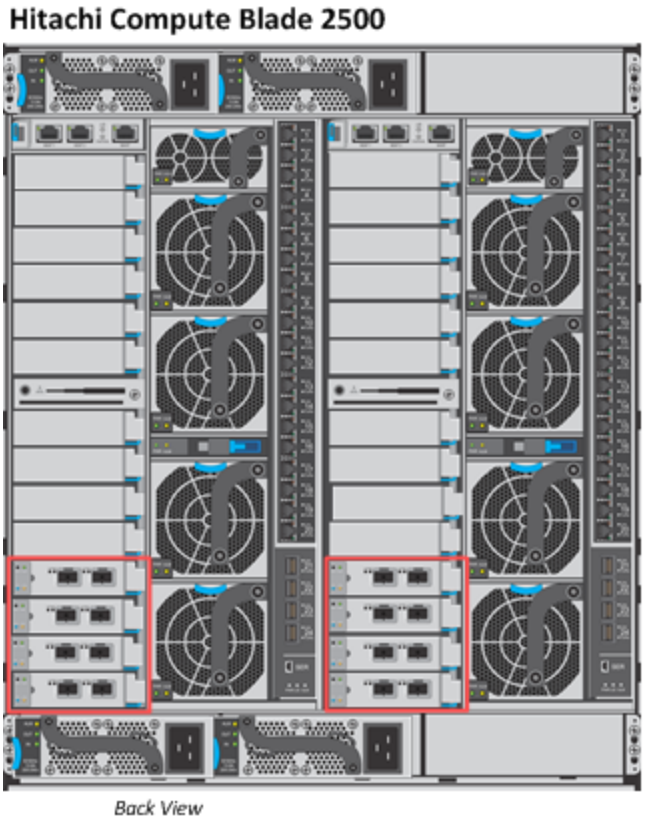


**Figure 4**

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

**Table 10. Hitachi HBA PCIe Card Parameters**

| Setting | Parameter |
|---|---|
| Boot Function | Enable |
| Link Speed | 16Gbps |
| Connection Type | Point-to-Pont |
| Multiple Port ID | Disable |
| Select Boot Device | Enable |
| Multipath Function | Enable |

## Network Architecture

This architecture requires the following separate networks:

- **Private Network (also called cluster interconnect)** — This network must be scalable. In addition, it must meet the low latency needs of the network traffic generated by cache synchronization of Oracle RAC and inter-node communication amongst the nodes in the cluster.
- **Public Network** — This network provides client connections to the applications and Oracle RAC.

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

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

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

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

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

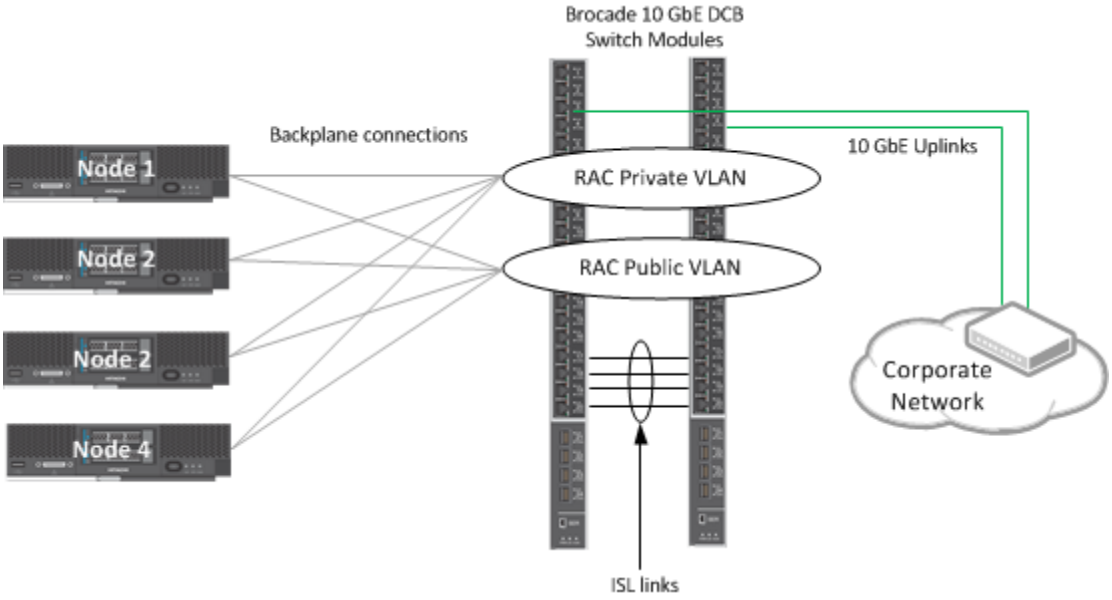


**Figure 5**

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

**Table 11. Network Configuration**

| Server | NIC Ports | UMC Physical Function Number | Switch Bay ID | Switch Ports (Internal) | VLAN | NIC BOND | Network | Bandwidth (Gb/sec) |
|---|---|---|---|---|---|---|---|---|
| Database Server 1 | B7-CNIC-0 | 0 | 1 | 1 | 1 | Bond1 | Private | 10 |
| | B7-CNIC-1 | 1 | 2 | 1 | | | | 10 |
| | B7-CNIC-2 | 2 | 1 | 15 | 2 | Bond2 | Public Oracle | 9 |
| | | 6 | 1 | 15 | 3 | Bond3 | Public Management | 1 |
| | B7-CNIC-3 | 3 | 2 | 15 | 2 | Bond2 | Public Oracle | 9 |
| | | 7 | 2 | 15 | 3 | Bond3 | Public Management | 1 |
| Database Server 2 | B8-CNIC-0 | 0 | 1 | 2 | 1 | Bond1 | Private | 10 |
| | B8-CNIC-1 | 1 | 2 | 2 | | | | 10 |
| | B8-CNIC-2 | 2 | 1 | 16 | 2 | Bond2 | Public Oracle | 9 |
| | | 6 | 1 | 16 | 3 | Bond3 | Public Management | 1 |
| | B8-CNIC-3 | 3 | 2 | 16 | 2 | Bond2 | Public Oracle | 9 |
| | | 7 | 2 | 16 | 3 | Bond3 | Public Management | 1 |

**Table 11. Network Configuration (Continued)**

| Server | NIC Ports | UMC Physical Function Number | Switch Bay ID | Switch Ports (Internal) | VLAN | NIC BOND | Network | Bandwidth (Gb/sec) |
|---|---|---|---|---|---|---|---|---|
| Database Server 3 | B9-CNIC-0 | 0 | 1 | 3 | 1 | Bond1 | Private | 10 |
| | B9-CNIC-1 | 1 | 2 | 3 | | | | 10 |
| | B9-CNIC-2 | 2 | 1 | 17 | 2 | Bond2 | Public Oracle | 9 |
| | | 6 | 1 | 17 | 3 | Bond3 | Public Management | 1 |
| | B9-CNIC-3 | 3 | 2 | 17 | 2 | Bond2 | Public Oracle | 9 |
| | | 7 | 2 | 17 | 3 | Bond3 | Public Management | 1 |
| Database Server 4 | B10-CNIC-0 | 0 | 1 | 4 | 1 | Bond1 | Private | 10 |
| | B10-CNIC-1 | 1 | 2 | 4 | | | | 10 |
| | B10-CNIC-2 | 2 | 1 | 18 | 2 | Bond2 | Public Oracle | 9 |
| | | 6 | 1 | 18 | 3 | Bond3 | Public Management | 1 |
| | B10-CNIC-3 | 3 | 2 | 18 | 2 | Bond2 | Public Oracle | 9 |
| | | 7 | 2 | 18 | 3 | Bond3 | Public Management | 1 |

# Test Results

This section summarizes the key observations from the test results from Hitachi Unified Compute Platform 6000 for Oracle Real Application Cluster using Hitachi Virtual Storage Platform F800 with Hitachi Accelerated Flash and Oracle Orion, Peakmarks, and SLOB.

Table 12 lists the Oracle Orion I/O test cases for this solution.

**Table 12. Oracle Orion I/O Test Case**

| Test Case | Metric | Value |
|---|---|---|
| 100% 8 KB Random Read | IOPS | 712,103 |
| | Latency | 1.1 ms |
| 75% 8 KB Random Reads + 25% 8 KB Random Write | IOPS | 300,657 |
| | Latency | 0.9 ms |
| 100% 1 MB Sequential Read | Throughput | 21,860 MB/sec |

Table 13 lists Peakmarks Storage test cases for this solution.

**Table 13. Peakmarks Storage Test Cases**

| Test Case | Metric | Value |
|---|---|---|
| STO-RR (100% 8KB Random Read) | IOPS | 529K |
| | Latency | 1.0 ms |
| STO-RW (100% 32 KB Random Write) | IOPS | 314K |
| | Throughput | 9,829 MB/sec |
| STO-SR (100% 1 MB Sequential Read) | Throughput | 24,335 MB/sec |
| Peakmarks DBX-U25 (Updating 25 Rows per Transaction) | Read IOPS | 118,240 |
| | Write IOPS | 126,807 |
| | Transactions per Second (TPS) | 4897.6 |

Table 14 lists SLOB test cases for this solution.

**Table 14. SLOB 25% Update OLTP Workload**

| Test Parameters/Test Results | VSP F800 | VSP F800 with Accelerated Compression On |
|---|---|---|
| SLOB Users | 116 | 96 |
| SLOB Tablespace Size (TB) | 9 | 9 |
| Storage Cache Read-hit Ratio | <1 % | <1% |
| DB File Sequential Read (ms) | 0.95 | 0.92 |
| Physical Read Total IO Requests per Second | 264,072 | 254,312 |
| Physical Write Total IO Requests per Second | 68,456 | 65,740 |
| Transactions per Second (TPS) | 1042.6 | 1010.5 |

# 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-542-01. September 2016.