TECHNICAL PAPER

HITACHI
Inspire the Next

Hitachi Unified Compute Platform for Oracle Database with the Oracle Real Application Clusters Option on 16 Nodes Using Hitachi Virtual Storage Platform G1000 with Hitachi Accelerated Flash

Reference Architecture Guide

By Kishore Daggubati, James Byun, and Debdatta Chattopadhyay

April 22, 2015

Hitachi Data Systems

Feedback

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

Table of Contents

Hitachi Unified Compute Platform for Oracle Database with the Oracle Real Application Clusters Option on 16 Nodes Using Hitachi Virtual Storage Platform G1000 with Hitachi Accelerated Flash

Reference Architecture Guide

This reference architecture guide shows how using the Hitachi Unified Compute Platform for Oracle Database with the Oracle Real Application Clusters (RAC) option provides a high performance, integrated solution for an Oracle infrastructure. The environment uses Hitachi Virtual Storage Platform G1000 (VSP G1000) with Hitachi Accelerated Flash. Use this resource to design an infrastructure for your requirements and budget.

This validated solution integrates servers, storage systems, and network and storage software. The environment provides reliability, high availability, scalability, and performance while processing small-scale to large-scale OLTP workloads. The dedicated server runs Oracle Database 12c R1 with the Oracle RAC option. The operating system is Oracle Linux using the Unbreakable Enterprise Kernel 6, Update 5.

You benefit from using this reference architecture if you are one of the following:

- Database administrator
- Storage administrator
- IT professional with the responsibility of planning and deploying an Oracle Database 12c R1 solution

To use this reference architecture guide, you need familiarity with the following:

- Hitachi Virtual Storage Platform G1000 using Hitachi Accelerated Flash
- Hitachi Compute Blade 2000
- Storage area networks
- Oracle RAC Database 12c Release 1
- Oracle Automatic Storage Management

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 for Oracle Database with the Real Application Clusters option using Virtual Storage Platform G1000 with Hitachi Accelerated Flash. It includes the high availability, performance, and scalability benefits of this environment for typical on-line transaction processing (OLTP) workloads. Tailor your implementation of this solution to meet your specific needs.

This reference architecture includes the following:

- Two Hitachi Compute Blade 2000 platforms, each with eight server blades
  - Each server blade hosts one Oracle database server
- Hitachi Virtual Storage Platform G1000 with Hitachi Accelerated Flash
- SAN infrastructure

Figure 1 shows the infrastructure for this solution.

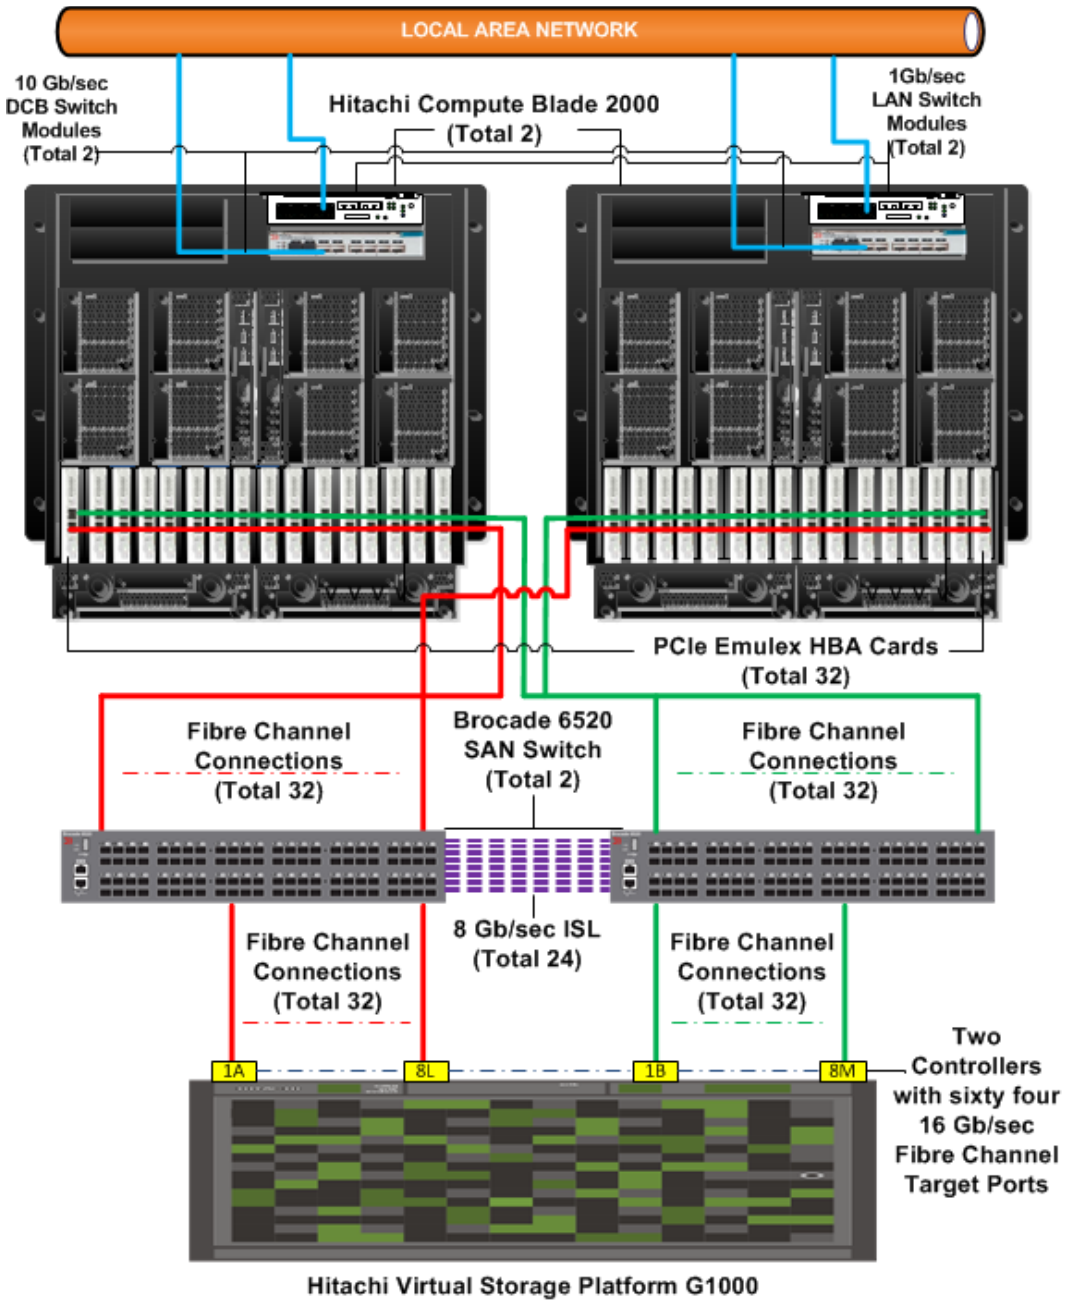

Figure 1

Key Solution Components

Table 1 and Table 2 list the key hardware and software components used in this reference architecture.

Table 1. Key Hardware Components

Component	Description	Version	Quantity
Server Chassis	Hitachi Compute Blade 2000	Firmware Version A0350-A-7248	2
Server Blade	Model GVAX55R3 (AX55 R3), each configured as follows: ■ Intel Xeon E5-2690 at 2.90 GHz, two 8-core physical processors ■ 64 GB RAM using 8 GB DIMMs ■ Two dual port 16 Gb/sec PCIe Fibre Channel cards ■ One onboard dual port 1 Gb/sec Ethernet NIC ■ One quad port 10 Gb/sec Ethernet NIC (Mezzanine Slot 0)	BMC Version 05-52 EFI Version 09-51	16

Storage System	Hitachi Virtual Storage Platform G1000 configured as follows:
	■ 2 flash chassis, each with:
	  ■ 32 × 1.6 TB flash modules in 4 flash units with 8 flash modules in each unit
	■ 102 TB raw capacity
	■ 2 main storage blades, each with:
	  ■ 4 front-end connectivity modules
	  ■ 8 × 16 Gb/sec Fibre Channel ports
	  ■ 4 back-end connectivity modules
	  ■ 8 × 6 Gb/sec SAS links
	■ 16 MP blades, each with two 8-core Intel XEON processors, 2.1 GHz
	■ 1 TB cache
	  ■ 16 GB shared memory size
	  ■ 10 GB cache directory size
	■ One front-end port on each controller connects to Hitachi Compute Blade 2000 over the SAN network
		80-01-22/01	1

Table 1. Key Hardware Components (Continued)

Component	Description	Version	Quantity
SAN Connectivity	■ Brocade 6520, 16 Gb/sec Fibre Channel switch	Firmware Version 7.1.0c	2
Network Connectivity	■ 1 Gb LAN switch module	Microcode dependent	1
	■ 10 Gb DCB switch module	Microcode dependent	1

Note — This reference architecture uses two Brocade 6520, 16 Gb/sec, Fibre Channel switches for SAN connectivity. You can use specific Fibre Channel switches to meet your needs.

Table 2. Key Software Components

Component	Description	Version
Operating System	Oracle Linux using Unbreakable Enterprise Kernel	6 Update 5
Volume Manager and File System Software	Oracle Automatic Storage Management	12c R1, 12.1.0.1
Database Software	Oracle	12c R1, 12.1.0.1
Cluster Software	Oracle Real Application Cluster	12c R1, 12.1.0.1
Storage Management Software	Hitachi Device Manager	Microcode dependent
Load Generator Software	Benchware	8.6 Build 140801

Hitachi Compute Blade 2000

Hitachi Compute Blade 2000 is an enterprise-class blade server platform. It features the following:

- A balanced system architecture that eliminates bottlenecks in performance and throughput
- Configuration flexibility
- Eco-friendly power-saving capabilities
- Fast server failure recovery using a N+1 cold standby design that allows replacing failed servers within minutes

Hitachi Virtual Storage Platform G1000

Hitachi Virtual Storage Platform G1000 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 G1000 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.

Hitachi Device Manager

Hitachi Device Manager enables essential management and optimization functions. Using Java agents, Storage Navigator runs on most browsers. A command line interface is available.

Use Device Manager for the following:

- Pool creation and expansion
- LUN creation and expansion
- Online microcode updates and other system maintenance functions
- Performance metrics

You need Device Manager to take advantage of the full features of Hitachi Virtual Storage Platform G1000.

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 6520, 16Gb/sec Fibre Channel switch

Oracle Linux

Oracle Linux is an enterprise-class operating system built and tested to run Oracle hardware, databases, and middleware. It is fully compatible with the Red Hat Enterprise Linux kernel.

Oracle Database with the Real Application Cluster Option

Oracle Database is optimized for use with Oracle products. This solution uses Oracle Real Application Cluster and Oracle Automatic Storage Management, which are part of the grid infrastructure component in Oracle Database.

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

Solution Design

This section describes the reference architecture environment implementing a large-sized environment for Hitachi Unified Compute Platform for Oracle Database with the Oracle Real Application Clusters option using Hitachi Virtual Storage Platform G1000 with Hitachi Accelerated Flash.

Specific infrastructure configuration includes the following:

- **Server** — This is a single server node, consisting of one AX55 R3 server blade for one logical Oracle Database server.
- **Storage System** — There are LDEVs mapped to each port that are presented to the server as LUNs.
- **SAN Fabric** — There are thirty-two zones created on each Fibre Channel switch module to zone the two PCIe Fibre Channel dual ports on each server blade and the sixty-four storage host 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

Create the RAID groups with Hitachi Device Manager on Hitachi Virtual Storage Platform G1000 as shown in Figure 2.

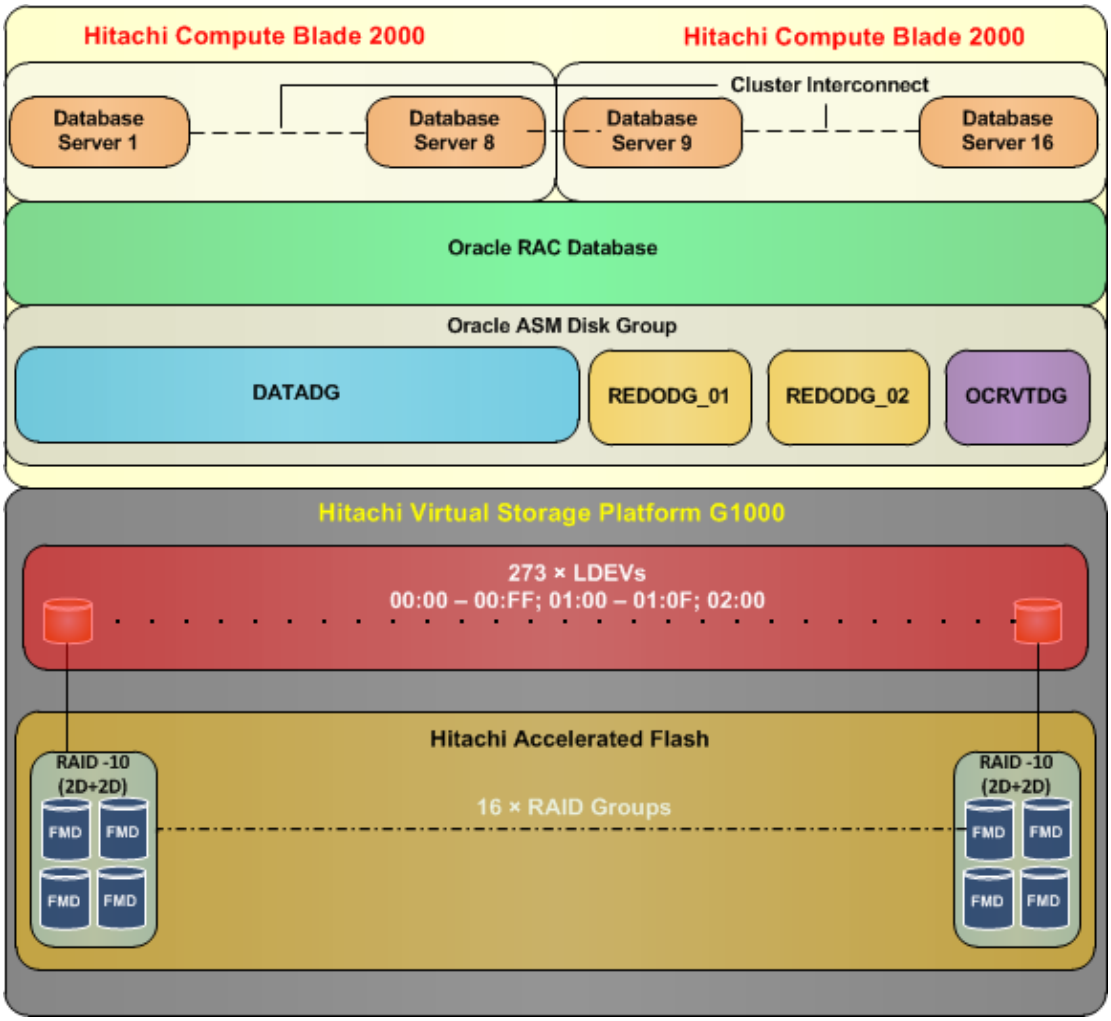

Figure 2

Table 3 has the details of the RAID groups for the Oracle RAC Database.

Table 3. RAID Groups for the Oracle RAC Database

RAID Group	*Purpose*	*RAID Level*	*Drive Type*	*No of Drives per RAID Group*	*Capacity (GB) per RAID Group*
1-1 - 1-4 2-1 - 2-4 3-1 - 3-4 4-1 - 4-4	Oracle RAC Database	RAID-10 (2D+2D)	1.6 TB Flash Module Drives	4	3200

Table 4 through Table 19 have the details for the LDEVs created in these RAID groups:

Table 4. LDEVs in Parity Group 1-1

Parity Group	LDEVs	LDEV Size (GB)	Purpose	Storage Port
1-1	00:00	150	Oracle RAC Database	1A,2A,1J,2J,1C,2C,1L 2L,3A,4A,3J,4J,3C,4C, 3L,4L
	00:10			
	00:20			
	00:30			
	00:40			
	00:50			
	00:60			
	00:70			
	00:80			
	00:90			
	00:A0			
	00:B0			
	00:C0			
	00:D0			
	00:E0			
	00:F0			
	01:00	60		
	02:00		Oracle Cluster Registry and Voting Disk	

Table 5. LDEVs in Parity Group 1-2

Parity Group	*LDEVs*	*LDEV Size (GB)*	*Purpose*	*Storage Port*
1-2	00:01	150	Oracle RAC Database	1A,2A,1J,2J,1C,2C,1L 2L,3A,4A,3J,4J,3C,4C, 3L,4L
	00:11			
	00:21			
	00:31			
	00:41			
	00:51			
	00:61			
	00:71			
	00:81			
	00:91			
	00:A1			
	00:B1			
	00:C1			
	00:D1			
	00:E1			
	00:F1			
	01:01	60		

Table 6. LDEVs in Parity Group 1-3

Parity Group	*LDEVs*	*LDEV Size (GB)*	*Purpose*	*Storage Port*
1-3	00:02	150	Oracle RAC Database	1A,2A,1J,2J,1C,2C,1L 2L,3A,4A,3J,4J,3C,4C, 3L,4L
	00:12			
	00:22			
	00:32			
	00:42			
	00:52			
	00:62			
	00:72			
	00:82			
	00:92			
	00:A2			
	00:B2			
	00:C2			
	00:D2			
	00:E2			
	00:F2			
	01:02	60		

Table 7. LDEVs in Parity Group 1-4

Parity Group	*LDEVs*	*LDEV Size (GB)*	*Purpose*	*Storage Port*
1-4	00:03	150	Oracle RAC Database	1A,2A,1J,2J,1C,2C,1L 2L,3A,4A,3J,4J,3C,4C, 3L,4L
	00:13			
	00:23			
	00:33			
	00:43			
	00:53			
	00:63			
	00:73			
	00:83			
	00:93			
	00:A3			
	00:B3			
	00:C3			
	00:D3			
	00:E3			
	00:F3			
	01:03	60		

Table 8. LDEVs in Parity Group 2-1

Parity Group	*LDEVs*	*LDEV Size (GB)*	*Purpose*	*Storage Port*
2-1	00:04	150	Oracle RAC Database	5A,6A,5J,6J,5C,6C,5L, 6L,7A,8A,7J,8J,7C,8C, 7L,8L
	00:14			
	00:24			
	00:34			
	00:44			
	00:54			
	00:64			
	00:74			
	00:84			
	00:94			
	00:A4			
	00:B4			
	00:C4			
	00:D4			
	00:E4			
	00:F4			
	01:04	60		

Table 9. LDEVs in Parity Group 2-2

Parity Group	*LDEVs*	*LDEV Size (GB)*	*Purpose*	*Storage Port*
2-2	00:05	150	Oracle RAC Database	5A,6A,5J,6J,5C,6C,5L, 6L,7A,8A,7J,8J,7C,8C, 7L,8L
	00:15			
	00:25			
	00:35			
	00:45			
	00:55			
	00:65			
	00:75			
	00:85			
	00:95			
	00:A5			
	00:B5			
	00:C5			
	00:D5			
	00:E5			
	00:F5			
	01:05	60		

Table 10. LDEVs in Parity Group 2-3

Parity Group	*LDEVs*	*LDEV Size (GB)*	*Purpose*	*Storage Port*
2-3	00:06	150	Oracle RAC Database	5A,6A,5J,6J,5C,6C,5L, 6L,7A,8A,7J,8J,7C,8C, 7L,8L
	00:16			
	00:26			
	00:36			
	00:46			
	00:56			
	00:66			
	00:76			
	00:86			
	00:96			
	00:A6			
	00:B6			
	00:C6			
	00:D6			
	00:E6			
	00:F6			
	01:06	60		

Table 11. LDEVs in Parity Group 2-4

Parity Group	LDEVs	LDEV Size (GB)	Purpose	Storage Port
2-4	00:07	150	Oracle RAC Database	5A,6A,5J,6J,5C,6C,5L, 6L,7A,8A,7J,8J,7C,8C, 7L,8L
	00:17			
	00:27			
	00:37			
	00:47			
	00:57			
	00:67			
	00:77			
	00:87			
	00:97			
	00:A7			
	00:B7			
	00:C7			
	00:D7			
	00:E7			
	00:F7			
	01:07	60		

Table 12. LDEVs in Parity Group 3-1

Parity Group	*LDEVs*	*LDEV Size (GB)*	*Purpose*	*Storage Port*
3-1	00:08	150	Oracle RAC Database	1B,2B,1K,2K,1D,2D,1M,2M,3B,4B,3K,4K,3D,4D,3M,4M
	00:18			
	00:28			
	00:38			
	00:48			
	00:58			
	00:68			
	00:78			
	00:88			
	00:98			
	00:A8			
	00:B8			
	00:C8			
	00:D8			
	00:E8			
	00:F8			
	01:08	60		

Table 13. LDEVs in Parity Group 3-2

Parity Group	*LDEVs*	*LDEV Size (GB)*	*Purpose*	*Storage Port*
3-2	00:09	150	Oracle RAC Database	1B,2B,1K,2K,1D,2D,1M,2M,3B,4B,3K,4K,3D,4D,3M,4M
	00:19			
	00:29			
	00:39			
	00:49			
	00:59			
	00:69			
	00:79			
	00:89			
	00:99			
	00:A9			
	00:B9			
	00:C9			
	00:D9			
	00:E9			
	00:F9			
	01:09	60		

Table 14. LDEVs in Parity Group 3-3

Parity Group	*LDEVs*	*LDEV Size (GB)*	*Purpose*	*Storage Port*
3-3	00:0A	150	Oracle RAC Database	1B,2B,1K,2K,1D,2D,1M,2M,3B,4B,3K,4K,3D,4D,3M,4M
	00:1A			
	00:2A			
	00:3A			
	00:4A			
	00:5A			
	00:6A			
	00:7A			
	00:8A			
	00:9A			
	00:AA			
	00:BA			
	00:CA			
	00:DA			
	00:EA			
	00:FA			
	01:0A	60		

Table 15. LDEVs in Parity Group 3-4

Parity Group	*LDEVs*	*LDEV Size (GB)*	*Purpose*	*Storage Port*
3-4	00:0B	150	Oracle RAC Database	1B,2B,1K,2K,1D,2D,1M,2M,3B,4B,3K,4K,3D,4D,3M,4M
	00:1B			
	00:2B			
	00:3B			
	00:4B			
	00:5B			
	00:6B			
	00:7B			
	00:8B			
	00:9B			
	00:AB			
	00:BB			
	00:CB			
	00:DB			
	00:EB			
	00:FB			
	01:0B	60		

Table 16. LDEVs in Parity Group 4-1

Parity Group	*LDEVs*	*LDEV Size (GB)*	*Purpose*	*Storage Port*
4-1	00:0C	150	Oracle RAC Database	5B,6B,5K,6K,5D,6D,5M,6M,7B,8B,7K,8K,8D,8D,7M,8M
	00:1C			
	00:2C			
	00:3C			
	00:4C			
	00:5C			
	00:6C			
	00:7C			
	00:8C			
	00:9C			
	00:AC			
	00:BC			
	00:CC			
	00:DC			
	00:EC			
	00:FC			
	01:0C	60		

Table 17. LDEVs in Parity Group 4-2

Parity Group	*LDEVs*	*LDEV Size (GB)*	*Purpose*	*Storage Port*
4-2	00:0D	150	Oracle RAC Database	5B,6B,5K,6K,5D,6D,5M,6M,7B,8B,7K,8K,8D,8D,7M,8M
	00:1D			
	00:2D			
	00:3D			
	00:4D			
	00:5D			
	00:6D			
	00:7D			
	00:8D			
	00:9D			
	00:AD			
	00:BD			
	00:CD			
	00:DD			
	00:ED			
	00:FD			
	01:0D	60		

Table 18. LDEVs in Parity Group 4-3

Parity Group	*LDEVs*	*LDEV Size (GB)*	*Purpose*	*Storage Port*
4-3	00:0E	150	Oracle RAC Database	5B,6B,5K,6K,5D,6D,5M,6M,7B,8B,7K,8K,8D,8D,7M,8M
	00:1E			
	00:2E			
	00:3E			
	00:4E			
	00:5E			
	00:6E			
	00:7E			
	00:8E			
	00:9E			
	00:AE			
	00:BE			
	00:CE			
	00:DE			
	00:EE			
	00:FE			
	01:0E	60		

Table 19. LDEVs in Parity Group 4-4

Parity Group	LDEVs	LDEV Size (GB)	Purpose	Storage Port
4-4	00:0F	150	Oracle RAC Database	5B,6B,5K,6K,5D,6D,5M,6M,7B,8B,7K,8K,8D,8D,7M,8M
	00:1F			
	00:2F			
	00:3F			
	00:4F			
	00:5F			
	00:6F			
	00:7F			
	00:8F			
	00:9F			
	00:AF			
	00:BF			
	00:CF			
	00:DF			
	00:EF			
	00:FF			
	01:0F	60		

Database Layout

The database layout design uses recommended practices from Hitachi Data Systems for Hitachi Virtual Storage Platform G1000 using Hitachi Accelerated Flash for small random I/O traffic, such as OLTP transactions. The layout also takes into account the Oracle ASM best practices when using Hitachi storage.

Base the storage design for database layout needs on the requirements of a specific application implementation. The design can vary greatly from one implementation to another. The components in this solution set have the flexibility for use in various deployment scenarios to provide the right balance between performance and ease of management for a given scenario.

- **Data and Indexes Tablespace** — Assign *Data ASM diskgroup* for the data and index tablespaces. The small file table space consists of 655 data files that are each 8 GB. Set the tablespace to a small initial size with auto extend enabled to maximize storage utilization.
- **Temp Tablespace** — Place *TEMP tablespace* in this configuration in the Data ASM diskgroup. Quite a number of small file tempfiles are created within one single small TEMP tablespace. Limit the size of each small file tempfile to 31.25 GB.
- **Undo Tablespace** — Place *UNDO tablespace* in this configuration in the Data ASM diskgroup. Assign one UNDO tablespace for each database instance in a sixteen node Oracle RAC database. Quite a number of small file undo datafiles are created within each small UNDO tablespace. Limit the size of each small undo datafile to 8 GB.
- **Online Redo Logs** — Assign one ASM diskgroup for each database instance. Four redo logs are created for each database instance in a sixteen node Oracle RAC database. Set the size of each redo log file to 8 GB.
- **Oracle Cluster Registry and Voting Disk** — Place each of these files in this configuration in the Archive ASM diskgroup.
- **Size Settings** — Set the database block size to 8 KB. Set the ASM allocation unit to 4 MB.
- **Memory Settings** — Set the Oracle database memory settings as follows:
  - SGA_TARGET = 32 GB.
  - PGA_AGGREGATE_TARGET = 8 GB
  - DB_CACHE_SIZE = 8 GB
  - DB_KEEP_CACHE_SIZE = 8 GB
  - DB_RECYCLE_CACHE_SIZE = 1 GB
- **ASM FILE SYSTEM I/O Settings** — Set the Oracle ASM I/O operations for database files as follows:
  - FILESYSTEMIO_OPTIONS = setall

Table 20 has the details for the disk mappings from the LUNS to the operating system devices and to the ASM disk groups for Oracle RAC Database tablespaces.

Table 20. Oracle ASM Disk Mapping for Oracle RAC Database Tablespace Details

ASM Disk Group	ASM Disk	Partition	LDEV Size (GB)	Parity Group	LDEV ID	MPU
DATADG	ASMDATA01	/dev/sda1	150	1-1	00:00	0
	ASMDATA02	/dev/sdb1	150	1-2	00:01	
	ASMDATA03	/dev/sdc1	150	1-3	00:02	
	ASMDATA04	/dev/sdd1	150	1-4	00:03	
	ASMDATA05	/dev/sde1	150	2-1	00:04	
	ASMDATA06	/dev/sdf1	150	2-2	00:05	
	ASMDATA07	/dev/sdg1	150	2-2	00:06	
	ASMDATA08	/dev/sdh1	150	2-4	00:07	
	ASMDATA09	/dev/sdi1	150	3-1	00:08	
	ASMDATA10	/dev/sdj1	150	3-2	00:09	
	ASMDATA11	/dev/sdk1	150	3-3	00:0A	
	ASMDATA12	/dev/sdl1	150	3-4	00:0B	
	ASMDATA13	/dev/sdm1	150	4-1	00:0C	
	ASMDATA14	/dev/sdn1	150	4-2	00:0D	
	ASMDATA15	/dev/sdo1	150	4-3	00:0E	
	ASMDATA16	/dev/sdp1	150	4-4	00:0F	
	ASMDATA17	/dev/sdq1	150	1-1	00:10	1
	ASMDATA18	/dev/sdr1	150	1-2	00:11	
	ASMDATA19	/dev/sds1	150	1-3	00:12	
	ASMDATA20	/dev/sdt1	150	1-4	00:13	
	ASMDATA21	/dev/sdu1	150	2-1	00:14	
	ASMDATA22	/dev/sdv1	150	2-2	00:15	
	ASMDATA23	/dev/sdw1	150	2-3	00:16	
	ASMDATA24	/dev/sdx1	150	2-4	00:17	
	ASMDATA25	/dev/sdy1	150	3-1	00:18	
	ASMDATA26	/dev/sdz1	150	3-2	00:19	
	ASMDATA27	/dev/sdaa1	150	3-3	00:1A	
	ASMDATA28	/dev/sdab1	150	3-4	00:1B	
	ASMDATA29	/dev/sdac1	150	4-1	00:1C	
	ASMDATA30	/dev/sdad1	150	4-2	00:1D	
	ASMDATA31	/dev/sdae1	150	4-3	00:1E	
	ASMDATA32	/dev/sdaf1	150	4-4	00:1F	

Table 20. Oracle ASM Disk Mapping for Oracle RAC Database Tablespace Details (Continued)

ASM Disk Group	ASM Disk	Partition	LDEV Size (GB)	Parity Group	LDEV ID	MPU
	ASMDATA33	/dev/sdag1	150	1-1	00:20	2
	ASMDATA34	/dev/sdah1	150	1-2	00:21	
	ASMDATA35	/dev/sdai1	150	1-3	00:22	
	ASMDATA36	/dev/sdaj1	150	1-4	00:23	
	ASMDATA37	/dev/sdak1	150	2-1	00:24	
	ASMDATA38	/dev/sdal1	150	2-2	00:25	
	ASMDATA39	/dev/sdam1	150	2-3	00:26	
	ASMDATA40	/dev/sdan1	150	2-4	00:27	
	ASMDATA41	/dev/sdao1	150	3-1	00:28	
	ASMDATA42	/dev/sdap1	150	3-2	00:29	
	ASMDATA43	/dev/sdaq1	150	3-3	00:2A	
	ASMDATA44	/dev/sdar1	150	3-4	00:2B	
	ASMDATA45	/dev/sdas1	150	4-1	00:2C	
	ASMDATA46	/dev/sdat1	150	4-2	00:2D	
	ASMDATA47	/dev/sdau1	150	4-3	00:2E	
	ASMDATA48	/dev/sdav1	150	4-4	00:2F	
	ASMDATA49	/dev/sdaw1	150	1-1	00:30	3
	ASMDATA50	/dev/sdax1	150	1-2	00:31	
	ASMDATA51	/dev/sday1	150	1-3	00:32	
	ASMDATA52	/dev/sdaz1	150	1-4	00:33	
	ASMDATA53	/dev/sdba1	150	2-1	00:34	
	ASMDATA54	/dev/sdbb1	150	2-2	00:35	
	ASMDATA55	/dev/sdbc1	150	2-3	00:36	
	ASMDATA56	/dev/sdbd1	150	2-4	00:37	
	ASMDATA57	/dev/sdbe1	150	3-1	00:38	
	ASMDATA58	/dev/sdbf1	150	3-2	00:39	
	ASMDATA59	/dev/sdbg1	150	3-3	00:3A	
	ASMDATA60	/dev/sdbh1	150	3-4	00:3B	
	ASMDATA61	/dev/sdbi1	150	4-1	00:3C	
	ASMDATA62	/dev/sdbj1	150	4-2	00:3D	
	ASMDATA63	/dev/sdbk1	150	4-3	00:3E	
	ASMDATA64	/dev/sdbl1	150	4-4	00:3F	

Table 20. Oracle ASM Disk Mapping for Oracle RAC Database Tablespace Details (Continued)

ASM Disk Group	ASM Disk	Partition	LDEV Size (GB)	Parity Group	LDEV ID	MPU
	ASMDATA65	/dev/sdbm1	150	1-1	00:40	4
	ASMDATA66	/dev/sdbn1	150	1-2	00:41	
	ASMDATA67	/dev/sdbo1	150	1-3	00:42	
	ASMDATA68	/dev/sdbp1	150	1-4	00:43	
	ASMDATA69	/dev/sdbq1	150	2-1	00:44	
	ASMDATA70	/dev/sdbr1	150	2-2	00:45	
	ASMDATA71	/dev/sdbs1	150	2-2	00:46	
	ASMDATA72	/dev/sdbt1	150	2-4	00:47	
	ASMDATA73	/dev/sdbu1	150	3-1	00:48	
	ASMDATA74	/dev/sdbv1	150	3-2	00:49	
	ASMDATA75	/dev/sdbw1	150	3-3	00:4A	
	ASMDATA76	/dev/sdbx1	150	3-4	00:4B	
	ASMDATA77	/dev/sdby1	150	4-1	00:4C	
	ASMDATA78	/dev/sdbz1	150	4-2	00:4D	
	ASMDATA79	/dev/sdca1	150	4-3	00:4E	
	ASMDATA80	/dev/sdcb1	150	4-4	00:4F	
	ASMDATA81	/dev/sdcc1	150	1-1	00:50	5
	ASMDATA82	/dev/sdcd1	150	1-2	00:51	
	ASMDATA83	/dev/sdce1	150	1-3	00:52	
	ASMDATA84	/dev/sdcf1	150	1-4	00:53	
	ASMDATA85	/dev/sdcg1	150	2-1	00:54	
	ASMDATA86	/dev/sdch1	150	2-2	00:55	
	ASMDATA87	/dev/sdci1	150	2-3	00:56	
	ASMDATA88	/dev/sdcj1	150	2-4	00:57	
	ASMDATA89	/dev/sdck1	150	3-1	00:58	
	ASMDATA90	/dev/sdcl1	150	3-2	00:59	
	ASMDATA91	/dev/sdcm1	150	3-3	00:5A	
	ASMDATA92	/dev/sdcn1	150	3-4	00:5B	
	ASMDATA93	/dev/sdco1	150	4-1	00:5C	
	ASMDATA94	/dev/sdcp1	150	4-2	00:5D	
	ASMDATA95	/dev/sdcq1	150	4-3	00:5E	
	ASMDATA96	/dev/sdcr1	150	4-4	00:5F	

Table 20. Oracle ASM Disk Mapping for Oracle RAC Database Tablespace Details (Continued)

ASM Disk Group	ASM Disk	Partition	LDEV Size (GB)	Parity Group	LDEV ID	MPU
	ASMDATA97	/dev/sdcs1	150	1-1	00:60	6
	ASMDATA98	/dev/sdct1	150	1-2	00:61	
	ASMDATA98	/dev/sdcu1	150	1-3	00:62	
	ASMDATA100	/dev/sdcv1	150	1-4	00:63	
	ASMDATA101	/dev/sdcw1	150	2-1	00:64	
	ASMDATA102	/dev/sdcx1	150	2-2	00:65	
	ASMDATA103	/dev/sdcy1	150	2-3	00:66	
	ASMDATA104	/dev/sdcz1	150	2-4	00:67	
	ASMDATA105	/dev/sdda1	150	3-1	00:68	
	ASMDATA106	/dev/sddb1	150	3-2	00:69	
	ASMDATA107	/dev/sddc1	150	3-3	00:6A	
	ASMDATA108	/dev/sddd1	150	3-4	00:6B	
	ASMDATA109	/dev/sdde1	150	4-1	00:6C	
	ASMDATA110	/dev/sddf1	150	4-2	00:6D	
	ASMDATA111	/dev/sddg1	150	4-3	00:6E	
	ASMDATA112	/dev/sddh1	150	4-4	00:6F	
	ASMDATA113	/dev/sddi1	150	1-1	00:70	7
	ASMDATA114	/dev/sddj1	150	1-2	00:71	
	ASMDATA115	/dev/sddk1	150	1-3	00:72	
	ASMDATA116	/dev/sddl1	150	1-4	00:73	
	ASMDATA117	/dev/sddm1	150	2-1	00:74	
	ASMDATA118	/dev/sddn1	150	2-2	00:75	
	ASMDATA119	/dev/sddo1	150	2-3	00:76	
	ASMDATA120	/dev/sddp1	150	2-4	00:77	
	ASMDATA121	/dev/sddq1	150	3-1	00:78	
	ASMDATA122	/dev/sddr1	150	3-2	00:79	
	ASMDATA123	/dev/sdds1	150	3-3	00:7A	
	ASMDATA124	/dev/sddt1	150	3-4	00:7B	
	ASMDATA125	/dev/sddu1	150	4-1	00:7C	
	ASMDATA126	/dev/sddv1	150	4-2	00:7D	
	ASMDATA127	/dev/sddw1	150	4-3	00:7E	
	ASMDATA128	/dev/sddx1	150	4-4	00:7F	

Table 20. Oracle ASM Disk Mapping for Oracle RAC Database Tablespace Details (Continued)

ASM Disk Group	ASM Disk	Partition	LDEV Size (GB)	Parity Group	LDEV ID	MPU
	ASMDATA129	/dev/sddy1	150	1-1	00:80	8
	ASMDATA130	/dev/sddz1	150	1-2	00:81	
	ASMDATA131	/dev/sdea1	150	1-3	00:82	
	ASMDATA132	/dev/sdeb1	150	1-4	00:83	
	ASMDATA133	/dev/sdec1	150	2-1	00:84	
	ASMDATA134	/dev/sded1	150	2-2	00:85	
	ASMDATA135	/dev/sdee1	150	2-2	00:86	
	ASMDATA136	/dev/sdef1	150	2-4	00:87	
	ASMDATA137	/dev/sdeg1	150	3-1	00:88	
	ASMDATA138	/dev/sdeh1	150	3-2	00:89	
	ASMDATA139	/dev/sdei1	150	3-3	00:8A	
	ASMDATA140	/dev/sdej1	150	3-4	00:8B	
	ASMDATA141	/dev/sdek1	150	4-1	00:8C	
	ASMDATA142	/dev/sdel1	150	4-2	00:8D	
	ASMDATA143	/dev/sdem1	150	4-3	00:8E	
	ASMDATA144	/dev/sden1	150	4-4	00:8F	
	ASMDATA145	/dev/sdeo1	150	1-1	00:90	9
	ASMDATA146	/dev/sdep1	150	1-2	00:91	
	ASMDATA147	/dev/sdeq1	150	1-3	00:92	
	ASMDATA148	/dev/sder1	150	1-4	00:93	
	ASMDATA149	/dev/sdes1	150	2-1	00:94	
	ASMDATA150	/dev/sdet1	150	2-2	00:95	
	ASMDATA151	/dev/sdeu1	150	2-3	00:96	
	ASMDATA152	/dev/sdev1	150	2-4	00:97	
	ASMDATA153	/dev/sdew1	150	3-1	00:98	
	ASMDATA154	/dev/sdex1	150	3-2	00:99	
	ASMDATA155	/dev/sdey1	150	3-3	00:9A	
	ASMDATA156	/dev/sdez1	150	3-4	00:9B	
	ASMDATA157	/dev/sdfa1	150	4-1	00:9C	
	ASMDATA158	/dev/sdfb1	150	4-2	00:9D	
	ASMDATA159	/dev/sdfc1	150	4-3	00:9E	
	ASMDATA160	/dev/sdfd1	150	4-4	00:9F	

Table 20. Oracle ASM Disk Mapping for Oracle RAC Database Tablespace Details (Continued)

ASM Disk Group	ASM Disk	Partition	LDEV Size (GB)	Parity Group	LDEV ID	MPU
	ASMDATA161	/dev/sdfe1	150	1-1	00:A0	10
	ASMDATA162	/dev/sdff1	150	1-2	00:A1	
	ASMDATA163	/dev/sdfg1	150	1-3	00:A2	
	ASMDATA164	/dev/sdfh1	150	1-4	00:A3	
	ASMDATA165	/dev/sdfi1	150	2-1	00:A4	
	ASMDATA166	/dev/sdfj1	150	2-2	00:A5	
	ASMDATA167	/dev/sdfk1	150	2-3	00:A6	
	ASMDATA168	/dev/sdfl1	150	2-4	00:A7	
	ASMDATA169	/dev/sdfm1	150	3-1	00:A8	
	ASMDATA170	/dev/sdfn1	150	3-2	00:A9	
	ASMDATA171	/dev/sdfo1	150	3-3	00:AA	
	ASMDATA172	/dev/sdfp1	150	3-4	00:AB	
	ASMDATA173	/dev/sdfq1	150	4-1	00:AC	
	ASMDATA174	/dev/sdfr1	150	4-2	00:AD	
	ASMDATA175	/dev/sdfs1	150	4-3	00:AE	
	ASMDATA176	/dev/sdft1	150	4-4	00:AF	
	ASMDATA177	/dev/sdfu1	150	1-1	00:B0	11
	ASMDATA178	/dev/sdfv1	150	1-2	00:B1	
	ASMDATA179	/dev/sdfw1	150	1-3	00:B2	
	ASMDATA180	/dev/sdfx1	150	1-4	00:B3	
	ASMDATA181	/dev/sdfy1	150	2-1	00:B4	
	ASMDATA182	/dev/sdfz1	150	2-2	00:B5	
	ASMDATA183	/dev/sdga1	150	2-3	00:B6	
	ASMDATA184	/dev/sdgb1	150	2-4	00:B7	
	ASMDATA185	/dev/sdgc1	150	3-1	00:B8	
	ASMDATA186	/dev/sdgd1	150	3-2	00:B9	
	ASMDATA187	/dev/sdge1	150	3-3	00:BA	
	ASMDATA188	/dev/sdgf1	150	3-4	00:BB	
	ASMDATA189	/dev/sdgg1	150	4-1	00:BC	
	ASMDATA190	/dev/sdgh1	150	4-2	00:BD	
	ASMDATA191	/dev/sdgi1	150	4-3	00:BE	
	ASMDATA192	/dev/sdgj1	150	4-4	00:BF	

Table 20. Oracle ASM Disk Mapping for Oracle RAC Database Tablespace Details (Continued)

ASM Disk Group	ASM Disk	Partition	LDEV Size (GB)	Parity Group	LDEV ID	MPU
	ASMDATA193	/dev/sdgk1	150	1-1	00:C0	12
	ASMDATA194	/dev/sdgl1	150	1-2	00:C1	
	ASMDATA195	/dev/sdgm1	150	1-3	00:C2	
	ASMDATA196	/dev/sdgn1	150	1-4	00:C3	
	ASMDATA197	/dev/sdgo1	150	2-1	00:C4	
	ASMDATA198	/dev/sdgp1	150	2-2	00:C5	
	ASMDATA199	/dev/sdgq1	150	2-2	00:C6	
	ASMDATA200	/dev/sdgr1	150	2-4	00:C7	
	ASMDATA201	/dev/sdgs1	150	3-1	00:C8	
	ASMDATA202	/dev/sdgt1	150	3-2	00:C9	
	ASMDATA203	/dev/sdgu1	150	3-3	00:CA	
	ASMDATA204	/dev/sdgv1	150	3-4	00:CB	
	ASMDATA205	/dev/sdgw1	150	4-1	00:CC	
	ASMDATA206	/dev/sdgx1	150	4-2	00:CD	
	ASMDATA207	/dev/sdgy1	150	4-3	00:CE	
	ASMDATA208	/dev/sdgz1	150	4-4	00:CF	
	ASMDATA209	/dev/sdha1	150	1-1	00:D0	13
	ASMDATA210	/dev/sdhb1	150	1-2	00:D1	
	ASMDATA211	/dev/sdhc1	150	1-3	00:D2	
	ASMDATA212	/dev/sdhd1	150	1-4	00:D3	
	ASMDATA213	/dev/sdhe1	150	2-1	00:D4	
	ASMDATA214	/dev/sdhf1	150	2-2	00:D5	
	ASMDATA215	/dev/sdhg1	150	2-3	00:D6	
	ASMDATA216	/dev/sdhh1	150	2-4	00:D7	
	ASMDATA217	/dev/sdhi1	150	3-1	00:D8	
	ASMDATA218	/dev/sdhj1	150	3-2	00:D9	
	ASMDATA219	/dev/sdhk1	150	3-3	00:DA	
	ASMDATA220	/dev/sdhl1	150	3-4	00:DB	
	ASMDATA221	/dev/sdhm1	150	4-1	00:DC	
	ASMDATA222	/dev/sdhn1	150	4-2	00:DD	
	ASMDATA223	/dev/sdho1	150	4-3	00:DE	
	ASMDATA224	/dev/sdhp1	150	4-4	00:DF	

Table 20. Oracle ASM Disk Mapping for Oracle RAC Database Tablespace Details (Continued)

ASM Disk Group	ASM Disk	Partition	LDEV Size (GB)	Parity Group	LDEV ID	MPU
	ASMDATA225	/dev/sdhq1	150	1-1	00:E0	14
	ASMDATA226	/dev/sdhr1	150	1-2	00:E1	
	ASMDATA227	/dev/sdhs1	150	1-3	00:E2	
	ASMDATA228	/dev/sdht1	150	1-4	00:E3	
	ASMDATA229	/dev/sdhu1	150	2-1	00:E4	
	ASMDATA230	/dev/sdhv1	150	2-2	00:E5	
	ASMDATA231	/dev/sdhw1	150	2-3	00:E6	
	ASMDATA232	/dev/sdhx1	150	2-4	00:E7	
	ASMDATA233	/dev/sdhy1	150	3-1	00:E8	
	ASMDATA234	/dev/sdhz1	150	3-2	00:E9	
	ASMDATA235	/dev/sdia1	150	3-3	00:EA	
	ASMDATA236	/dev/sdib1	150	3-4	00:EB	
	ASMDATA237	/dev/sdic1	150	4-1	00:EC	
	ASMDATA238	/dev/sdid1	150	4-2	00:ED	
	ASMDATA239	/dev/sdie1	150	4-3	00:EE	
	ASMDATA240	/dev/sdif1	150	4-4	00:EF	15
	ASMDATA241	/dev/sdig1	150	1-1	00:F0	
	ASMDATA242	/dev/sdih1	150	1-2	00:F1	
	ASMDATA243	/dev/sdii1	150	1-3	00:F2	
	ASMDATA244	/dev/sdij1	150	1-4	00:F3	
	ASMDATA245	/dev/sdik1	150	2-1	00:F4	
	ASMDATA246	/dev/sdil1	150	2-2	00:F5	
	ASMDATA247	/dev/sdim1	150	2-3	00:F6	
	ASMDATA248	/dev/sdin1	150	2-4	00:F7	
	ASMDATA249	/dev/sdio1	150	3-1	00:F8	
	ASMDATA250	/dev/sdip1	150	3-2	00:F9	
	ASMDATA251	/dev/sdiq1	150	3-3	00:FA	
	ASMDATA252	/dev/sdir1	150	3-4	00:FB	
	ASMDATA253	/dev/sdis1	150	4-1	00:FC	
	ASMDATA254	/dev/sdit1	150	4-2	00:3FD	
	ASMDATA255	/dev/sdiu1	150	4-3	00:FE	
	ASMDATA256	/dev/sdiv1	150	4-4	00:FF	

Table 21 has the details for the disk mappings from the LUNS to the operating system devices and to the ASM disk groups for the Oracle online redo logs.

Table 21. Oracle ASM Disk Mapping for Oracle Redo Log Details

ASM Disk Group	ASM Disk	Partition	LDEV Size (GB)	Parity Group	LDEV ID	MPU
REDODG_01	ASMREDO01	/dev/sdiw1	60	1-1	01:00	0
	ASMREDO02	/dev/sdix1	60	1-2	01:01	1
	ASMREDO03	/dev/sdiy1	60	1-3	01:02	2
	ASMREDO04	/dev/sdiz1	60	1-4	01:03	3
	ASMREDO05	/dev/sdja1	60	2-1	01:04	4
	ASMREDO06	/dev/sdjb1	60	2-2	01:05	5
	ASMREDO07	/dev/sdjc1	60	2-3	01:06	6
	ASMREDO08	/dev/sdjd1	60	2-4	01:07	7
REDODG_02	ASMREDO09	/dev/sdje1	60	3-1	01:08	8
	ASMREDO10	/dev/sdjf1	60	3-2	01:09	9
	ASMREDO11	/dev/sdjg1	60	3-3	01:0A	10
	ASMREDO12	/dev/sdjh1	60	3-4	01:0B	11
	ASMREDO13	/dev/sdji1	60	4-1	01:0C	12
	ASMREDO14	/dev/sdjj1	60	4-2	01:0D	13
	ASMREDO15	/dev/sdjk1	60	4-3	01:0E	14
	ASMREDO16	/dev/sdjl1	60	4-4	01:0F	15

Table 22 has the details for the disk mappings from the LUNS to the operating system devices and to the ASM disk groups for the Oracle Cluster Registry and Voting disk.

Table 22. Oracle ASM Disk Mapping for Oracle Cluster Registry and Voting Disk Details

ASM Disk Group	ASM Disk	Partition	LDEV Size (GB)	Parity Group	LDEV ID	MPU
OCRVTDG	ASMOCR01	/dev/sdjm1	60	1-1	02:00	1

Server and Application Architecture

The reference architecture uses two Hitachi Compute Blade 2000 chassis. Each chassis includes eight server blades. Each server blade hosts one Oracle database server. The Oracle database server contains a single logical database instance.

There are 16 combined CPU cores and 32 combined CPU threads with 64 GB RAM for each database server. This provides the compute power for Oracle RAC Database to handle complex database queries and a large volume of transaction processing in parallel. See “Key Solution Components” on page 5 for hardware details.

Set the following parameters for Oracle Linux Unbreakable Enterprise Kernel 6, Update 5 as follows:

- I/O Scheduler = NOOP
- Single partition on each LUN with offset = 4 MB

Set the following parameter for each of the PCIe HBA Emulex cards as follows:

- Queue Depth = 128

Table 3 shows the server infrastructure for the reference architecture.

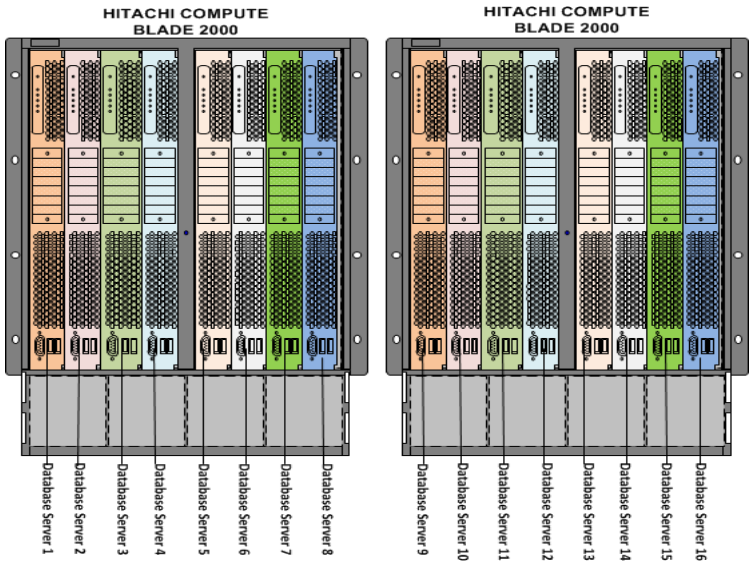

Figure 3

SAN Architecture

Map the provisioned LDEVs to multiple ports on Hitachi Virtual Storage Platform G1000 using Hitachi Accelerated Flash. These LDEV port assignments provide multiple paths to the storage system from the host for high availability.

The environment uses two Brocade 6520 Fibre Channel switch modules for scalability and high availability. “Key Solution Components” on page 5 has host configuration details.

Each of the database servers uses four Fibre Channel ports, with two ports from each of the PCIe HBA Emulex cards. This provides a sixty four path connection for all LUNs mapped to each of the database servers in a sixteen node Oracle RAC Database.

The environment uses two Brocade 6520 Fibre Channel switch modules to provide redundant paths for all Hitachi Virtual Storage Platform G1000 systems using Hitachi Accelerated Flash LUNs.

Note — This reference architecture used sixty four 16 Gb/sec Fibre Channel target ports and thirty two 16 Gb/sec PCIe HBA Emulex cards. Each of the target ports and HBA ports were configured with a speed of 8 Gb/sec. You may choose to use a speed of 16 Gb/sec to meet your needs.

Table 23 provides the zoning details for the SAN.

Table 23. SAN Zoning Details

Server	HBA Ports	Switch Zone	Storage Port	Switch
Database Server 1	CB50_HBA00	SW_01_CB50_HBA00	1A	BS118
	CB50_HBA01	SW_02_CB50_HBA01	1B	BS119
	CB50_HBA10	SW_01_CB50_HBA10	5A	BS118
	CB50_HBA11	SW_02_CB50_HBA11	5B	BS119
Database Server 2	CB51_HBA00	SW_01_CB51_HBA00	2A	BS118
	CB51_HBA01	SW_02_CB51_HBA01	2B	BS119
	CB51_HBA10	SW_01_CB51_HBA10	6A	BS118
	CB51_HBA11	SW_02_CB51_HBA11	6B	BS119
Database Server 3	CB52_HBA00	SW_01_CB52_HBA00	1J	BS118
	CB52_HBA01	SW_02_CB52_HBA01	1K	BS119
	CB52_HBA10	SW_01_CB52_HBA10	5J	BS118
	CB52_HBA11	SW_02_CB52_HBA11	5K	BS119
Database Server 4	CB53_HBA00	SW_01_CB53_HBA00	2J	BS118
	CB53_HBA01	SW_02_CB53_HBA01	2K	BS119
	CB53_HBA10	SW_01_CB53_HBA10	6J	BS118
	CB53_HBA11	SW_02_CB53_HBA11	6K	BS119
Database Server 5	CB54_HBA00	SW_01_CB54_HBA00	1C	BS118
	CB54_HBA01	SW_02_CB54_HBA01	1D	BS119
	CB54_HBA10	SW_01_CB54_HBA10	5C	BS118
	CB54_HBA11	SW_02_CB54_HBA11	5D	BS119
Database Server 6	CB55_HBA00	SW_01_CB55_HBA00	2C	BS118
	CB55_HBA01	SW_02_CB55_HBA01	2D	BS119
	CB55_HBA10	SW_01_CB55_HBA10	6C	BS118
	CB55_HBA11	SW_02_CB55_HBA11	6D	BS119
Database Server 7	CB56_HBA00	SW_01_CB56_HBA00	1L	BS118
	CB56_HBA01	SW_02_CB56_HBA01	5L	BS119
	CB56_HBA10	SW_01_CB56_HBA10	1M	BS118
	CB56_HBA11	SW_02_CB56_HBA11	5M	BS119

Table 23. SAN Zoning Details (Continued)

Server	HBA Ports	Switch Zone	Storage Port	Switch
Database Server 8	CB57_HBA00	SW_01_CB57_HBA00	2L	BS118
	CB57_HBA01	SW_02_CB57_HBA01	2M	BS119
	CB57_HBA10	SW_01_CB57_HBA10	6L	BS118
	CB57_HBA11	SW_02_CB57_HBA11	6M	BS119
Database Server 9	CB58_HBA00	SW_01_CB58_HBA00	3A	BS118
	CB58_HBA01	SW_02_CB58_HBA01	3B	BS119
	CB58_HBA10	SW_01_CB58_HBA10	7A	BS118
	CB58_HBA11	SW_02_CB58_HBA11	7B	BS119
Database Server 10	CB59_HBA00	SW_01_CB59_HBA00	4A	BS118
	CB59_HBA01	SW_02_CB59_HBA01	4B	BS119
	CB59_HBA10	SW_01_CB59_HBA10	8A	BS118
	CB59_HBA11	SW_02_CB59_HBA11	8B	BS119
Database Server 11	CB60_HBA00	SW_01_CB60_HBA00	3J	BS118
	CB60_HBA01	SW_02_CB60_HBA01	3K	BS119
	CB60_HBA10	SW_01_CB60_HBA10	7J	BS118
	CB60_HBA11	SW_02_CB60_HBA11	7K	BS119
Database Server 12	CB61_HBA00	SW_01_CB61_HBA00	4J	BS118
	CB61_HBA01	SW_02_CB61_HBA01	4K	BS119
	CB61_HBA10	SW_01_CB61_HBA10	8J	BS118
	CB61_HBA11	SW_02_CB61_HBA11	8K	BS119
Database Server 13	CB62_HBA00	SW_01_CB62_HBA00	3C	BS118
	CB62_HBA01	SW_02_CB62_HBA01	3D	BS119
	CB62_HBA10	SW_01_CB62_HBA10	7C	BS118
	CB62_HBA11	SW_02_CB62_HBA11	7D	BS119
Database Server 14	CB63_HBA00	SW_01_CB63_HBA00	4C	BS118
	CB63_HBA01	SW_02_CB63_HBA01	4D	BS119
	CB63_HBA10	SW_01_CB63_HBA10	8C	BS118
	CB63_HBA11	SW_02_CB63_HBA11	8D	BS119
Database Server 15	CB64_HBA00	SW_01_CB64_HBA00	3L	BS118
	CB64_HBA01	SW_02_CB64_HBA01	3M	BS119
	CB64_HBA10	SW_01_CB64_HBA10	7L	BS118
	CB64_HBA11	SW_02_CB64_HBA11	7M	BS119
Database Server 16	CB65_HBA00	SW_01_CB65_HBA00	4L	BS118
	CB65_HBA01	SW_02_CB65_HBA01	4M	BS119
	CB65_HBA10	SW_01_CB65_HBA10	8L	BS118
	CB65_HBA11	SW_02_CB65_HBA11	8M	BS119

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 Real Application Clusters and inter-node communication amongst the nodes in the cluster.
- **Public Network** — This network provides client connections to the applications and Oracle Real Application Clusters.

Hitachi Data Systems recommends using a pair of 1/10 Gb/sec NICs for the cluster interconnect and public network.

Each server in this reference architecture has one dual port 1 Gb/sec on-board NIC and one quad port 10 Gb/sec Ethernet NIC (Mezzanine Slot 0) for different types of data traffic. The NIC ports on each server blade connect to the internal 1/10 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 4 shows the network configuration for the reference architecture environment.

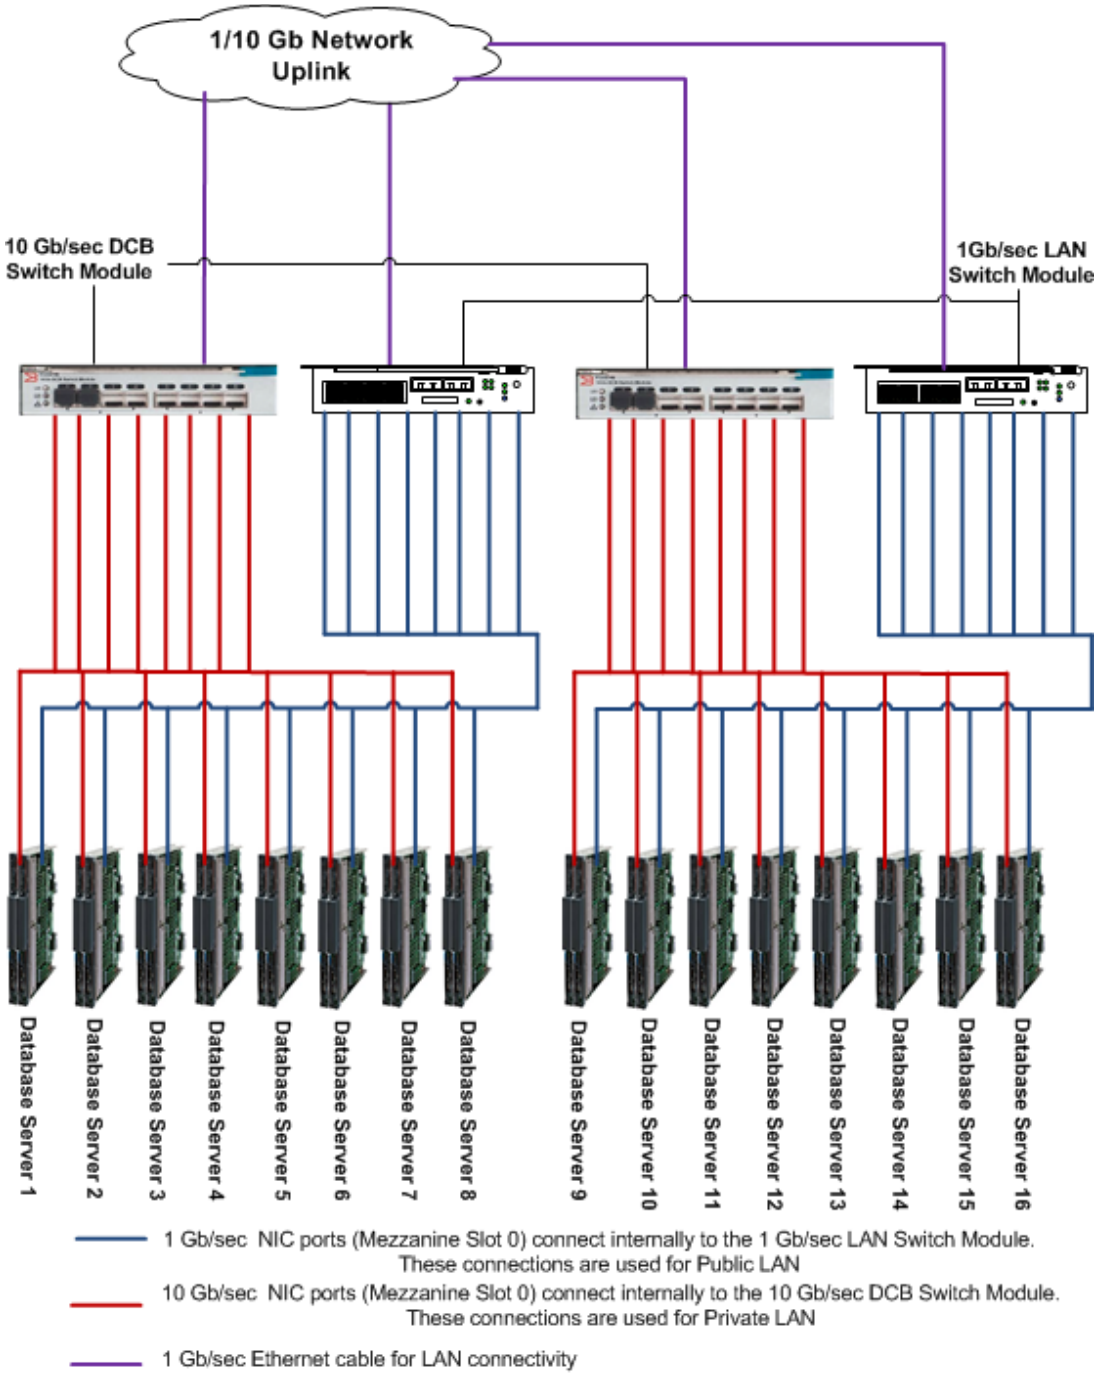

Figure 4

Note — This reference architecture uses one 1 Gb/sec LAN switch module and one 10 Gb/sec DCB switch module. For redundancy and throughput, you can use additional LAN switches to meet your needs.

Table 24 provides the network details for this solution.

Table 24. Network Details

Server	NIC Ports	Switch Bay ID	Switch Ports (Internal)	VLAN	Subnet	IP Address	Network
Database server 1	B0-NIC-0	0	5	17	17	172.17.17.50	Public
	B0-MIC-1	1	5	25	1	169.254.25.147	Private
Database server 2	B1-NIC-0	0	6	17	17	172.17.17.51	Public
	B1-MIC-1	1	6	25	1	169.254.173.71	Private
Database server 3	B2-NIC-0	0	7	17	17	172.17.17.52	Public
	B2-MIC-1	1	7	25	1	169.254.90.214	Private
Database server 4	B3-NIC-0	0	8	17	17	172.17.17.53	Public
	B3-MIC-1	1	8	25	1	169.254.164.179	Private
Database server 5	B4-NIC-0	0	9	17	17	172.17.17.54	Public
	B4-MIC-1	1	9	25	1	169.254.129.12	Private
Database server 6	B5-NIC-0	0	10	17	17	172.17.17.55	Public
	B5-MIC-1	1	10	25	1	169.254.50.246	Private
Database server 7	B6-NIC-0	0	11	17	17	172.17.17.56	Public
	B6-MIC-1	1	11	25	1	169.254.74.183	Private
Database server 8	B7-NIC-0	0	12	17	17	172.17.17.57	Public
	B7-MIC-1	1	12	25	1	169.254.90.75	Private
Database server 9	B8-NIC-0	0	13	17	17	172.17.17.58	Public
	B8-MIC-1	1	13	25	1	169.254.238.135	Private
Database server 10	B9-NIC-0	0	14	17	17	172.17.17.59	Public
	B9-MIC-1	1	14	25	1	169.254.101.100	Private
Database server 11	B10-NIC-0	0	15	17	17	172.17.17.60	Public
	B10-MIC-1	1	15	25	1	169.254.55.109	Private
Database server 12	B11-NIC-0	0	16	17	17	172.17.17.61	Public
	B11-MIC-1	1	16	25	1	169.254.162.229	Private
Database server 13	B12-NIC-0	0	17	17	17	172.17.17.62	Public
	B12-MIC-1	1	17	25	1	169.254.251.214	Private
Database server 14	B13-NIC-0	0	18	17	17	172.17.17.63	Public
	B13-MIC-1	1	18	25	1	169.254.67.141	Private
Database server 15	B14-NIC-0	0	19	17	17	172.17.17.64	Public
	B14-MIC-1	1	19	25	1	169.254.175.26	Private
Database server 16	B15-NIC-0	0	8	17	17	172.17.17.65	Public
	B15-MIC-1	1	8	25	1	169.254.57.206	Private

Engineering Validation

This section describes the validation of this large-sized Hitachi Unified Compute Platform for Oracle Database with the Oracle RAC option using Hitachi Virtual Storage Platform G1000 with Hitachi Accelerated Flash solution in this reference architecture for functionality and performance. Testing involved configuring a sixteen node Oracle RAC Database.

Test Methodology

This is the methodology used to validate this solution.

Workload

Testing included test iterations to run simulated and synthetic workloads using Benchware.

These were the simulated workloads for Oracle Benchware testing:

- Large sequential reads
- Small random reads
- I/O Latency for small random reads
- Oracle data load for OLTP
- SQL SELECT and UPDATE transactions

Data Gathering

Performance statistics were collected at the following levels:

- Storage
  - Hitachi Device Manager was used to collect storage performance data
- Operating System
  - “iostat” and “vmstat” were used to collect operating system statistics
- Database
  - Oracle Automatic Workload Repository report was used for database performance
  - Benchware was used for application-level statistics for the number of executed transactions
- Standalone tests on the drives/LUNS
  - Benchware was used to test the throughput and IOPS performance for a given workload

Test Results

This section summarizes the key observations from the test results for Hitachi Unified Compute Platform for Oracle Database with the Oracle RAC option and using Hitachi Virtual Storage Platform G1000 with Hitachi Accelerated Flash and Oracle Benchware.

Table 25 lists the throughput for the large sequential I/O test.

Table 25. Throughput for Large Sequential I/O

Test Case	*Throughput in MB/sec*
Sequential Reads	26,934

Table 26 lists the IOPS for the small random reads test.

Table 26. IOPS for Small Random Reads

Test Case	*IOPS*
Random Reads	3,179,157

Table 27 lists the latency for the small random reads test.

Table 27. Latency for Small Random Reads

Test Case	*Latency in ms*
Random Reads	1.08

Table 28 lists the throughput for Oracle uncompressed data load for OLTP.

Table 28. Oracle Data Load for OLTP

Test Case	*Throughput in Rows per second (rps)*
Conventional Load	623,900
Bulk Load with Eight Node Oracle RAC Database	5,697,000

Table 29 lists the throughput for Oracle SQL transactions with an eight node Oracle RAC Database.

Table 29. Throughput for Oracle SQL Transactions with an Eight Node Oracle RAC Database

Test Case	*Throughput in Transactions per second (tps)*
SELECT	1,153,000
UPDATE	378,400

Table 30 lists the latency for Oracle SQL transactions with an eight node Oracle RAC Database.

Table 30. Latency for Oracle SQL Transactions with an Eight Node Oracle RAC Database

Test Case	*Latency in ms*
SELECT	0.59
UPDATE	1.25

Conclusion

Testing this reference architecture with Benchware validated that Hitachi Unified Compute Platform for Oracle Database with the Oracle Real Application Clusters option can achieve up to the following:

- 3,179,157 random Oracle reads
- 26.30 GB/sec sequential Oracle reads

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 Corporation

Corporate Headquarters
2845 Lafayette Street, Santa Clara, California 95050-2627 USA
www.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-387-00, April 2015