

Deploy Hitachi Unified Compute Platform Select for Oracle Database using Oracle Database 11g R2 Enterprise Edition in a Medium-sized Solution

Implementation Guide

By Kishore Daggubati

January 30, 2013



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

Tested Solution Components.....	2
Hitachi Compute Blade 2000.....	6
Hitachi Unified Storage VM.....	7
Hitachi Dynamic Provisioning.....	7
Hitachi Storage Navigator.....	8
Oracle Linux.....	8
Oracle Database	8
Fusion-io ioDrive2	8
PuTTY.....	9
Solution Implementation.....	10
Preparation.....	11
Configure Hitachi Compute Blade 2000.....	12
Configure Hitachi Unified Storage VM.....	17
Set Up a SAN Boot for the Operating System.....	26
Install Oracle Enterprise Linux 5, Update 5, on Oracle Database Server.....	28
Install and Configure Oracle Enterprise Linux 5, Update 5, Device Mapper.....	29
Install and Configure Fusion-io ioDrive2 PCI-e Flash Card Acceleration (Optional).....	29
Install Oracle Database 11g Release 2 Automatic Storage Management	31
Configure Oracle Database 11g Release 2 Automatic Storage Management.....	31
Install Oracle Database 11g Release 2	34
Create the Oracle Database	35

Deploy Hitachi Unified Compute Platform Select for Oracle Database using Oracle Database 11g R2 Enterprise Edition in a Medium-sized Solution

Implementation Guide

This implementation instructs you on how to deploy the medium-sized Hitachi Unified Compute Platform Select for Oracle Database environment, with or without PCI-e flash acceleration. Using servers and storage from Hitachi, this solution hosts Oracle Database 11g Release 2.

A description of the architecture in this paper is in [Deploy Hitachi Unified Compute Platform Select for Oracle Database using Oracle Database 11g R2 Enterprise Edition in a Medium-sized Solution Reference Architecture Guide \(AS-198-00\)](#).

To benefit from this implementation guide, you need to be a database administrator, storage administrator, or have the responsibility to plan and deploy an Oracle Database 11g Release 2 solution. You need familiarity with the following:

- Hitachi Unified Storage VM
- Hitachi Compute Blade 2000
- Storage area networks
- Oracle Database 11g Release 2
- Oracle Automatic Storage Management
- Fusion-io ioDrive2 flash storage cards

Note — These procedures were developed in a lab environment. Many things affect production environments beyond prediction or duplication in a lab environment. Follow recommended practice by conducting proof-of-concept testing for acceptable results before implementing this solution in your production environment. Test the implementation in a non-production, isolated test environment that otherwise matches your production environment.

Tested Solution Components

Figure 1 shows the infrastructure without PCI-e flash acceleration.




Figure 1

Figure 2 shows the infrastructure using PCI-e flash acceleration.




Figure 2

Table 1 lists the hardware components used in this implementation guide.

Table 1. Solution Hardware Components

<i>Component</i>	<i>Description</i>	<i>Version</i>	<i>Quantity</i>
Server chassis	Hitachi Compute Blade 2000	Firmware Version A0195-C-6443	1
Server blades	<p>Model GVAX57A1 (X57-A1), each configured as follows:</p> <ul style="list-style-type: none"> ■ Intel Xeon X7560 processors, 2.26 GHz, <ul style="list-style-type: none"> ■ 2 × 8-core physical processors ■ 96 GB RAM using 4 GB DIMMs ■ 1 dual port Fibre Channel card (Mezzanine Slot 0) ■ 2 × 1 Gb/sec Ethernet NICs ■ (Used only with architecture using PCI-e flash card) 2 Fusion-io ioDrive2 <ul style="list-style-type: none"> ■ 1.2 TB MLC ■ Version 3.1.5, Boot version 0.0.3.1 ■ In PCI-e expansion Slot 0 and PCI-e expansion Slot 1 	EFI BIOS Version 4.6.3.7	2

Table 1. Solution Hardware Components (Continued)

Component	Description	Version	Quantity
Storage system	<p>Hitachi Unified Storage VM, configured as follows:</p> <ul style="list-style-type: none"> ■ 144 × 600 GB SAS 10k RPM drives in 6 trays with 24 disks in each tray, for about 86 TB raw capacity ■ 2 main storage blades, each with: <ul style="list-style-type: none"> ■ 2 front-end connectivity modules <ul style="list-style-type: none"> ■ 4 × 8 Gb/sec Fibre Channel ports, 8 ports total ■ 2 back-end connectivity modules <ul style="list-style-type: none"> ■ 8 × 6 Gb/sec SAS links each, 16 links total ■ 2 MP blades, each with one 8-core Intel XEON processor, 2.33 GHz ■ 55 GB cache ■ One front-end port on each controller connects to Hitachi Compute Blade 2000. 	73-01-01-00/00	1
PCI-e flash cards <ul style="list-style-type: none"> ■ Only used when using flash acceleration 	Fusion-io ioDrive2 <ul style="list-style-type: none"> ■ 1.2 TB PCI-e flash cards 		4
SAN connectivity	<ul style="list-style-type: none"> ■ 6 port, 8 Gb/sec Fibre Channel switch modules 	V 642b	2
Symmetric multiprocessing (SMP)	<ul style="list-style-type: none"> ■ SMP connector creates one server from two GVAX57A1 server blades 	N/A	1

Note — This implementation guide uses X57-A1 (GVAX57A1) server blades with 4 GB DIMMs. You may upgrade to X57-A2 server blades using Intel Xeon Processor E7-8870. In addition, you may upgrade the 4 GB DIMMs to 8 GB DIMMs

Table 2 lists the software components used in this implementation guide.

Table 2. Solution Software Components

<i>Component</i>	<i>Description</i>	<i>Version</i>
Operating system	Oracle Enterprise Linux	5 Update 5, Red Hat-compatible kernel
Volume manager and file system software	Oracle Automatic Storage Management	11g R2, 11.2.0.1.0
Multipath software	Device Mapper Multipath (DM-Multipath) of Red Hat Enterprise Linux for the Red Hat Enterprise Linux 5 release	0.4.7-34.el5
Database software	Oracle	11g R2, 11.2.0.1.0
Storage management software	Hitachi Storage Navigator	73-01-01/00
	Hitachi Dynamic Provisioning	Microcode dependent
I/O calibration software	Orion	11.1
Load generator software	Swingbench	2.4.0.845
Database client communication software	Oracle Net Services	11g R2, 11.2.0.1.0

Note — This implementation guide uses Oracle Linux 5, Update 5, and Oracle Database 11g R2. Instead of using Oracle Linux 5, Update 5, you can use Red Hat Linux 5 or later. If you are using PCI-e flash acceleration, you must use Oracle Database 11g or later. If you are not using PCI-e flash acceleration, you can use Oracle Database 9i or later.

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 Unified Storage VM

[Hitachi Unified Storage VM](#) is an entry-level enterprise storage platform. It combines storage virtualization services with unified block, file, and object data management. This versatile, scalable platform offers a storage virtualization system to provide central storage services to existing storage assets.

Unified management delivers end-to-end central storage management of all virtualized internal and external storage on Unified Storage VM. A unique, hardware-accelerated, object-based file system supports intelligent file tiering and migration, as well as virtual NAS functionality, without compromising performance or scalability.

The benefits of Unified Storage VM are the following:

- Enables the move to a new storage platform with less effort and cost when compared to the industry average
- Increases performance and lowers operating cost with automated data placement
- Supports scalable management for growing and complex storage environments while using fewer resources
- Achieves better power efficiency and with more storage capacity for more sustainable data centers
- Lowers operational risk and data loss exposure with data resilience solutions
- Consolidates management with end-to-end virtualization to prevent virtual server sprawl

Hitachi Dynamic Provisioning

On Hitachi storage systems, [Hitachi Dynamic Provisioning](#) provides wide striping and thin provisioning functionalities.

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

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

When used with [Hitachi Unified Storage VM](#), Hitachi Dynamic Provisioning has the benefit of thin provisioning. Physical space assignment from the pool to the dynamic provisioning volume happens as needed using 42 MB pages, up to the logical size specified for each dynamic provisioning volume. There can be a dynamic expansion or reduction of pool capacity without disruption or downtime. You can rebalance an expanded pool across the current and newly added RAID groups for an even striping of the data and the workload.

Hitachi Storage Navigator

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

Use Storage Navigator for the following:

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

You need Storage Navigator to take advantage of the full features of Hitachi Unified Storage VM.

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

[Oracle Database](#) is optimized for use with 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.

Fusion-io ioDrive2

[Fusion-io ioDrive 2](#) combines VSL and ioMemory into an ioMemory platform on a PCI-e flash card. The ioMemory platform provides consistent low latency access for mixed workloads with 15-microsecond access latency, 1.5 GB/sec bandwidth, over 275,000 read IOPS (512 bytes) and over 800,000 write IOPS (512 bytes).

This solution uses Fusion-io ioDrive2 as part of the default configuration when using PCI-e flash acceleration. They are high performance block devices used to accelerate Oracle database input and output performance.

PuTTY

Use [PuTTY](#) terminal emulator when performing configuration tasks in this implementation guide. It is a free implementation of TELNET and Secure Shell (SSH) protocols for Microsoft Windows and UNIX platforms.

Solution Implementation

The following are the processes used to deploy a medium-sized environment for Hitachi Unified Compute Platform Select for Oracle Database using Hitachi Unified Storage VM and Hitachi Compute Blade 2000.

- “Preparation” on page 11
- “Configure Hitachi Compute Blade 2000” on page 12
- “Configure Hitachi Unified Storage VM” on page 17
- “Set Up a SAN Boot for the Operating System” on page 26
- “Install Oracle Enterprise Linux 5, Update 5, on Oracle Database Server” on page 28
- “Install and Configure Oracle Enterprise Linux 5, Update 5, Device Mapper” on page 29
- “Install and Configure Fusion-io ioDrive2 PCI-e Flash Card Acceleration (Optional)” on page 29
- “Install Oracle Database 11g Release 2 Automatic Storage Management ” on page 31
- “Configure Oracle Database 11g Release 2 Automatic Storage Management” on page 31
- “Install Oracle Database 11g Release 2 ” on page 34
- “Create the Oracle Database ” on page 35

Some of these processes require that you use the following Hitachi Compute Blade 2000 interfaces:

- **Management Module Web GUI** — Manage the entire chassis using a web-based GUI. Open a web browser and enter the IP address of the management module.
 - **Server blade Web GUI** — Manage a server blade using a web-based GUI. Open a web browser and enter the IP address of a server blade.
 - **Server blade command-line interface** — Configure the SMP using the command-line interface.
-

For more information, see the following documentation:

- Hitachi Storage Navigator on-line help
- Hitachi Dynamic Provisioning user guide
- Hitachi Compute Blade 2000 user guide
- Oracle 11g Release 2 database installation guide
- Oracle Enterprise Linux 5, Update 5, Device Mapper user guide
- Fusion-io ioDrive2 knowledge base

Preparation

Prepare the following before implementing this solution:

- Verify that you have the following items:
 - USB DVD drive (comes with Hitachi Compute Blade 2000)
 - Microsoft Windows Server 2008 installation DVD or ISO image
 - Driver kit version 12-01 CD or ISO image
 - Includes driver for mezzanine Emulex HBA
 - Seven IP addresses:
 - **Management module** — 2
 - **Server blades** — 2
 - **Blade management console (BMC)** — 2
 - **Operating system** — 1
 - Four IP addresses in environments where these items are not already in place and connected:
 - Two for the storage systems
 - Two for the Fibre Channel switches
 - Make the following LAN connections:
 - Connect one network cable from each network switch module to your corporate LAN. Do not use the first network switch module port because it is reserved for system settings.
 - Connect one cable from each management module to your corporate management LAN.
-

Configure Hitachi Compute Blade 2000

These procedures to configure Hitachi Compute Blade 2000 assume that the following has happened:

- Racking of the server blade chassis
- Installation of the server blades into the chassis
- Completion of all cabling (LAN and SAN)

Configure Management IP Addresses

To configure the IP addresses for the management module and the server blades using the management module Web graphical user interface, do the following.

1. Connect a system console (laptop or PC) to the MGMT0 port of the management module with an Ethernet cable.
 - If you have two management modules installed, connect to the management module with the solid green MSR LED lamp, as shown in Figure 3.




Figure 3

2. Open a browser and type `http://192.168.0.1/` in the Address bar.
A logon screen displays.
3. Log on using these default user credentials:
User name: administrator
Password: password
The Web graphical user interface launches.

4. Open the **Management LAN network** window:
 - (1) Click the **Settings** tab.
 - (2) In the navigation tree on the left, click **The configuration of network** link.
5. Configure the management LAN network:
 - (1) Click **Edit**.

The fields become editable.
 - (2) Configure the management LAN network by typing the following for the management module and each partition:
 - IP address
 - Subnet mask
 - Default gateway

You need a different IP address for the management module and each partition (server blade).
 - (3) To save the settings, click **Confirm**, and then click **Apply**.
6. Disconnect the Ethernet cable from the management module.
7. Connect an Ethernet cable from the management module to an external management LAN switch.
 - This connection allows you to perform management through any system console in the management LAN.

Configure the SMP Connection

Before configuring SMP, verify the connection of the two server blades using an SMP connector.

To configure the SMP connection for the two sever blades, do the following.

1. Open a PuTTY session, using the following information:
 - **IP address** — 192. 168. 0. 1
 - **Connection type** — Tel net
 - **User name** — admi ni strator
 - **Password** — password
 2. To enter system command mode, type the following: S
-

3. To change the physical partition setting, type the following;
 - (1) PR (physical partition setting)
 - (2) C (edit partition configuration)
 - (3) 0 (lowest number partition)
 - (4) 2 (number of blades in the SMP configuration)
 - (5) Y (confirm the deletion of Partition1)
4. To save the physical partition settings, type the following:
 - (1) W (write configuration)
 - (2) Y (confirm)

The configuration displays in the PuTTY session, as shown in Figure 4.

```

172.17.170.150 - PuTTY
  C. Edit partition configuration.
  Q. Quit.
(S,P,R,C,[Q]) : C

-- Partition configuration --
      Server blade
P Power      0 1 2 3 4 5 6 7
-----
      S=S - - - - -
-----
0 On          X X - - - - -
1 -----
2 -----      - X - - - - -
3 -----      - - X - - - -
4 -----      - - - X - - -
5 -----      - - - - X - -
6 -----      - - - - - X -
7 -----      - - - - - - X

-- Edit partition configuration. --
  C. Configure partition.
  W. Write configuration.
  Q. Quit.
(C,W,[Q]) : █

```

Figure 4

Configure LAN Switch Modules

Table 3 lists the default IP addresses for the switch modules.

Table 3. LAN Switch Module Settings

<i>Item</i>	<i>Default Setting</i>
Switch module 0	192.168.0.28
Switch module 1	192.168.0.29
Subnet Mask	255.255.255.0

The command line interface for the system supports two command input modes:

- Operation command mode
- Configuration command mode

Configuration of the LAN switch modules requires the use of both command input modes. Table 4 describes features of command input modes and details to access and exit command input modes.

Table 4. Command Input Modes

<i>Command Input Mode</i>	<i>User Level</i>	<i>Accessed by</i>	<i>Prompt</i>	<i>Exited by</i>	<i>Commands You Can Use</i>
Operation command	General user	Login: <user-ID>	>	>logout	Some operation commands
	System administrator	>enable	#	# disable	All operation commands
Configuration command	System administrator	# configure	(config)#	(config)# exit	All configuration commands

By default, the configuration of Switch 0 Port 1 is for dedicated VLAN management. It cannot communicate with a server blade and other external ports. If needed, reconfigure the other port to be Port 1.

Use a public network for the initial configuration.

To configure the LAN Switch Module 0 and Switch Module 1, do the following.

1. Connect the management PC to Switch 0 Port 1 with a LAN cable.
2. Connect to the default switch IP with TELNET, 192.168.0.28, using these defaults:

User name — admi ni strator

Password — password

3. Repeat these steps separately for Switch Module 0 and Switch Module 1.
 - (1) To enter the switch module, type the following: `SW`
 - (2) To select the switch module, type the number of the module (either 0 or 1),
 - (3) To confirm the selection of Switch Module 0 or Switch Module 1, type the following: `Y`
 - (4) To log on, type the following: `operator`
A password is not required.
 - (5) To change to administrator mode, type the following: `enable`
 - (6) To display the current running configuration of the switch module, type the following: `show run`
Every VLAN, switch port, and VLAN port displays.
 - (7) To enter configuration mode, type the following: `config`
The cursor changes to this: `(config)#`
 - (8) Create and name VLAN 100.
 - i. To create VLAN 100, type the following: `vlan 100`
 - ii. To assign a name to VLAN 100, type the following: `test vlan`
 - (9) Save the configuration changes.
 - i. To change back to config mode, type the following: `exit`
 - ii. To write the changes to the configuration, type the following: `save`
-

Configure Hitachi Unified Storage VM

Use these processes to configure Hitachi Unified Storage VM.

Note — You must have **modify** privileges when using Hitachi Storage Navigator to complete these processes.

Open Hitachi Storage Navigator

Initially, use PuTTY to connect open Hitachi Storage Navigator using the default SVP IP address, 172.17.45.190. Later you will use a Web browser.

To open Hitachi Storage Navigator, do the following.

1. Connect the management PC to LAN Port 1 on the storage array with a LAN cable.
2. Power on the SVP.
 - Remove the front bezel panel to power on the SVP.Storage Navigator is set up on the SVP and is therefore ready for access.
3. Open a web browser and type the following in the Address bar:

`https://172.17.45.190`

4. Use the following default user credentials to log on to Storage Navigator:

User ID — system

Password — manager

5. Verify all the settings are correct.

Configure the Storage Area Network

Hitachi Unified Compute Platform Select for Oracle Database, with and without PCI-e flash acceleration, uses Port 1A, 2A, 1B, and 2B on Hitachi Unified Storage VM.

Use Device Mapper in Oracle Enterprise Linux 5, Update 5, to set the **round-robin** load-balancing algorithm.

Zone SAN Switches

The Oracle database server has four Fibre Channel ports, with two ports from the mezzanine card on each server blade. These mezzanine cards connect internally to the Fibre Channel switch modules located in the Hitachi Compute Blade 2000 chassis.

Connect two ports from each of the two internal Fibre Channel switch modules to the four ports on the Hitachi Unified Storage VM storage system as shown in Figure 1 on page 2 and Figure 2 on page 3, This provides a four-path connection.

Configure each connection with zones on fabric switches to provide four zones for the Oracle database server.

In addition, follow these recommended practices:

- Use World Wide Port Name (WWPN) identification for all zoning configuration.
- Connect a minimum of two HBAs per server for multipath high availability.
- Disable all unused switch ports to increase security.
- Configure ports for point-to-point topology.
- Set ports to a specific speed. Do not use the auto negotiate setting.
- Use single initiator zoning.

To configure zoning for the SAN switches, do the following.

1. Create aliases for ports.
2. Create zones.
3. Add members (aliases) to the zones.
4. Save the zones.
5. Enable the zones.

Your organization's guidelines for zoning might vary. Check with your IT department.

Table 5 lists the database server zoning for this solution. Use the same zoning whether your implementation is with or without PCI-e flash acceleration.

Table 5. Database Server Zoning

<i>HBA Ports</i>	<i>Switch Zone</i>	<i>Storage Port</i>	<i>Switch</i>
B0-HBA1-1	BS2K_13_B0_HBA1_1_ASE45_190_0A	1A	5300-05
B0-HBA1-2	BS2K_13_B0_HBA1_2_ASE45_190_1B	1B	5300-06
B1-HBA1-1	BS2K_13_B1_HBA1_1_ASE45_190_2A	2A	5300-05
B1-HBA1-2	BS2K_13_B1_HBA1_2_ASE45_190_2B	2B	5300-06

Configure Storage for Oracle Database Server

The procedures that you use to configure the storage assume you have installed all appropriate licenses on your storage system before starting.

Configure Fibre Channel Port Settings

Perform this process for each Fibre Channel storage port.

To configure each Fibre Channel storage port using Hitachi Storage Navigator, do the following.

1. Log on to Storage Navigator.
2. Open the **Ports** tab and select the port row to edit.
 - (1) From the **Storage Systems** pane, click the **Ports/Host Groups** link.
 - (2) Click the **Ports** tab and select the port row.
 - (3) Click **Edit Ports**.
3. Make the following choices.
 - (1) From the **Port Attribute** list, click **Target**
 - (2) From the **Port Security** list, click **Enable**
 - (3) From the **Port Speed** list, click **Auto**
 - (4) From the **Fabric** list, click **ON**
 - (5) From the **Connection Type** list, click **P-to-P**
4. Click **Finish**.

The **Confirm** window opens.
5. Click **Confirm**, check all settings, and then click **Apply**.
 - If **Go to tasks window for status** is checked, the **Tasks** window opens.

Create RAID Groups

This solution uses four RAID groups for the following:

- Operating system boot and Oracle software
- Oracle Online redo logs (2 RAID groups)
- Oracle archived redo logs for this solution

All drives are the following: 600 GB 10k RPM SAS.

Table 6 on page 20 lists the RAID group configuration for each parity group.

Table 6. RAID Groups Assignment

<i>Parity Group</i>	<i>Purpose</i>	<i>RAID Level</i>	<i>No of Drives</i>	<i>Capacity (GB)</i>
2-19	Operating System Boot and Oracle Software	RAID-5 (3D+1P)	4	500
2-1	Oracle Online Redo Logs	RAID-10 (2D+2D)	4	1000
2-2	Oracle Online Redo Logs	RAID-10 (2D+2D)	4	1000
2-3	Oracle Archived Redo Logs	RAID-10 (2D+2D)	4	1000

To create a RAID group, use SVP.

Create LDEVs

Create four storage LDEVs for the operating system and Oracle database server, as shown in Table 7. Map all the storage LDEVs to storage ports 1A, 2A, 1B, and 2B.

Table 7. Operating System and Oracle Database Server LDEV Information

<i>Parity Group</i>	<i>LDEVs</i>	<i>LDEV Size (GB)</i>	<i>Purpose</i>
2-19	00:00:00	500	Operating System Boot and Oracle Software
2-1	00:00:01	1000	Oracle Online Redo Logs and Oracle Control File
2-2	00:00:02	1000	Oracle Online Redo Logs and Oracle Control File
2-3	00:00:03	1000	Oracle Archived Redo Logs

To create LDEVs using Hitachi Storage Navigator, do the following.

1. Log on to Storage Navigator.
2. Open the **Create LDEVs** window.
 - (1) From the **Storage Systems** pane, click the **Logical Devices** link.

The **Logical Devices** window opens.
 - (2) Click the **Create LDEVs** tab.

The **Create LDEVs** window opens.
3. Create the LDEVs.

Repeat these steps to create every LDEV in Table 7.

 - (1) Click the following choices from these lists:
 - From the **Provisioning Type** list, **Basic**
 - From the **Drive Type/RPM** list in the **Parity Group Selection** area, **SAS/10K**
 - From the **RAID Level** list in the **Parity Group Selection** area, click the value from Table 7

(2) Type the following values in these text boxes:

- The LDEV size in the **LDEV Capacity** box and then click **GB** from the list
- The number of LDEVs to create in the **Number of LDEVs** box
- The name in the **LDEV Name** box

(3) Click the following choices from these lists:

- From the **Format Type** list, **Quick Format**.
- From the **MP Unit ID** list, **Auto**

(4) Click **Add**.

This adds the created LDEVs to the **Selected LDEVs** table.

(5) Click **Finish**.

The **Confirm** window opens.

- To continue the operation for setting the LU path and defining a logical unit, click **Next**. For details about how to set the LU path, see [Create LU Paths](#).

(6) Click **Confirm**, check all settings, and then click **Apply**.

- If **Go to tasks window for status** is checked, the **Tasks** window opens.
-

Create Dynamic Provisioning Pools

Use one dynamic provisioning pool, as shown in Table 8.

Table 8. Dynamic Provisioning Pool

<i>Dynamic Provisioning Pool ID</i>	<i>Parity Group</i>	<i>RAID Level</i>	<i>Drive Type</i>	<i>No of Drives</i>	<i>Pool Capacity</i>
ora_datapool_01	1-1 – 1-12 1-21 – 1-30 2-11 – 2-18	RAID-10 (2D+2D)*2	600GB 10k RPM SAS	120	30 TB

To create a dynamic provisioning pool using Hitachi Storage Navigator software, do the following.

1. Log on to Storage Navigator.
2. Open the **Create Pools** window.
 - (1) To open the storage system, click the **HUS VM** link.
 - (2) In the **Storage System** pane, click the **Pool** heading.

The **Pool** window opens
 - (3) Click **Create Pools**.

The **Create Pools** window opens.
3. Create the dynamic provisioning pool
 - (1) From the **Pool Type** list, click **Dynamic Provisioning**.
 - (2) From the **Multi-Tier Pool** list, click **Disable**.
 - (3) From the **Pool Volume Selection** list, click **Manual**.
 - (4) From the **Drive Type/RPM** list, click the hard disk drive type and RPM.
 - (5) From the **RAID Level** list, click the RAID level.
 - (6) Click **Select Pool VOLs**.
 - (7) The **Select Pool VOLs** window opens.
 - In the **Available Pool Volumes** table, select the pool-VOL row to be associated with a pool, and then click **Add**
 - (8) This registers the selected pool-VOL into the **Selected Pool Volumes** table.
 - (9) Click **OK**.
 - (10) This applies the information in the **Selected Pool Volumes** table to **Total Selected Pool Volumes** and **Total Selected Capacity**.

4. Identify the dynamic provisioning pool.
 - (1) In the **Pool Name** text box, type the pool name.
 - (2) Click **Option**.
 - (3) In the **Initial Pool ID** text box, type the number of the initial pool ID (from 0 to 127).
 - (4) In the **Subscription Limit** text box, type an integer value from 0 to 65534 as the subscription rate (in percent) for the pool.
 - (5) In the **Warning Threshold** text box, type an integer value from 1 to 100 as the rate (in percent) for the pool.
 - The default value is 70%.
 - (6) In the **Depletion Threshold** text box, type an integer value from 1 to 100 as the rate (in percent) for the pool.
 - The default value is 80%.
 5. Click **Add**.
 6. The new pool is added to the **Selected Pools** table on the right.
 7. Click **Next**.
 8. The **Create LDEVs** window appears. Continue with **Create LDEVs** .
 9. In the **Confirm** window, click **Apply** to register the setting in the task.
 10. If the **Go to tasks window for status** check box is selected, the **Tasks** window opens.
-

Create DP-VOLs

Create 29 storage virtual volumes for the Oracle database server, as shown in either Table 9 (without PCI-e acceleration) or Table 10 (with PCI-e acceleration). Map all the storage virtual volumes to the storage ports 1A, 2A, 1B, and 2B.

Table 9. Dynamic Provisioning Pool Volume Information without PCI-e Acceleration

<i>Dynamic Provisioning Pool ID</i>	<i>LDEVs</i>	<i>LDEV Size (GB)</i>	<i>Purpose</i>
001	00:00:22 – 00:00:3E	200	<ul style="list-style-type: none"> ■ Oracle System ■ Sysaux ■ Undo ■ Temp ■ OLTP Application Tablespaces

Table 10. Dynamic Provisioning Pool Volume Information with PCI-e Acceleration

<i>Dynamic Provisioning Pool ID</i>	<i>LDEVs</i>	<i>LDEV Size (GB)</i>	<i>Purpose</i>
001	00:00:22 – 00:00:39	200	<ul style="list-style-type: none"> ■ 2nd Failure Group to host Application Tablespace
001	00:00:3A – 00:00:3E	200	<ul style="list-style-type: none"> ■ Oracle System ■ Sysaux ■ Undo ■ Temp

To create the LDEVs, dynamic provisioning pools, and DP-VOLs using Hitachi Storage Navigator, do the following.

1. Log on to Storage Navigator.
2. Create the DP-VOLs.

Repeat these steps to create every DP-VOL in Table 9 or Table 10.

- (1) From the **Storage Systems** pane, click the **Pools** link and then click the **Create LDEVs** tab.

The **Create LDEVs** window opens.

- (2) Set the provisioning parameters.
 - i. From the **Provisioning Type** list, click **Dynamic Provisioning**.

- ii. From the **Drive Type/RPM** list in the **Parity Group Selection** area, click **SAS/10K**.
 - iii. From the **RAID Level** list in **Parity Group Selection**, click **1(2D+2D)**.
- (3) Click **Select Pool**.
- The **Select Pool** window opens.
- (4) From the **Available Pools** list, click a pool and then click **OK**.
- (5) Set the parameters for the pool.
- In the **LDEV Capacity** box, type the capacity and then click **GB** from the list.
 - In **Number of LDEVs** box, type the number to be created.
 - In the **LDEV Name** box, type the name.
 - From the **Format Type** list, click **Quick Format**.
 - From the **MP Unit ID** list, Click **Auto**.
- (6) Click **Add**.
- This adds the created LDEVs to the **Selected LDEVs** table.
- (7) Click **Finish**.
- The **Confirm** window opens.
- (8) Click **Confirm**, check all settings, and then click **Apply**.
- If the **Go to tasks window for status** check box is selected, the **Tasks** window opens.

Create LU Paths

The Oracle database server has four Fibre Channel ports, with two ports from the mezzanine card on each server blade. These mezzanine cards connect internally to the Fibre Channel switch modules located in the Hitachi Compute Blade 2000 chassis.

Connect the two ports from each of the two internal Fibre Channel switch modules to the four ports on the Hitachi Unified Storage VM storage system, as shown in Figure 1 on page 2 and Figure 4 on page 14. This provides a four-path connection for all LUNs mapped to the server.

To create LU paths using Hitachi Storage Navigator software, do the following.

1. Log on to Storage Navigator.
2. Create every LU.

Repeat these steps for every LU in your solution.

- (1) From the **Storage Systems** pane, click the **Ports/Host Groups** link.
- (2) From the **Actions** menu, click **Logical Device**, and then click **Add LUN Paths**.
- (3) From the **Available LDEVs** list, select the LDEVs to be added, and then click **Add**.

The selected LDEVs are now in the **Selected LDEVs** list.

- (4) Click **Next**.
- (5) From the **Available Host Groups** list, click the host groups to be added, and then click **Add**.
- (6) Click **Next** and then click **Finish**.
- (7) Click **Confirm**, check all the settings, and then click **Apply**.
 - If the **Go to tasks window for status** check box is selected, the **Tasks** window opens.

Set Up a SAN Boot for the Operating System

SAN boot is using a volume on Hitachi Unified Storage VM to boot the operating system without using the internal disks of the server. Do the following to set up the SAN boot for the operating system.

Register the Boot Device

Use the server blade Web GUI to configure the HBAs for registering the boot device.

To configure HBAs on the Oracle database server, do the following.

1. Open a Web browser and type the IP address of the server blade in the Address bar.

A logon screen opens.
2. Click **Launch Remote Console** and log on to Remote Console.

User Name: user01

Password: pass01

Remote Console opens.

3. Change the Emulex BIOS.

- (1) From the **Keyboard** menu, click **Ctrl+Alt+Del** to start the system.
- (2) When “Press <Alt+E> or <Ctrl+E> to enter Emulex BIOS configuration Utility” displays during system startup, press Ctrl+E.
- (3) To change the BIOS, type the number shown on the LightPulse BIOS Utility screen.
- (4) Select **Configure Boot Devices**.
- (5) Type the boot entry number to register the boot device and then press Enter.

A list of boot devices displays.

- (6) Type the 2-digit number of the SAN boot device and then press Enter.
- (7) Type the 2-digit number (in hexadecimal) of the LUN of the SAN boot device and then press Enter.

The volumes display in the ascending order of LUNs, starting with volume of the SAN boot device.

- (8) Type 01 and then press Enter.

4. Check the registration of the boot device.

- (1) To boot the device, type 1 and then press Enter.

A screen with a list of saved boot devices displays.

- (2) Verify the device registration in the boot entry.

5. Reboot so the changes takes effect.

- (1) To close the BIOS configuration utility, type x.

- (2) To reboot, press Y.

Mount the Oracle Enterprise Linux 5, Update 5, ISO Image from a Remote CD/DVD Image

To perform this procedure, you need the Remote Console session from “Register the Boot Device.”

To mount the Oracle Enterprise Linux 5U5 ISO image from a Remote CD/DVD Image, take the following steps.

1. From the **Remote CD/DVD** menu, click **Redirect CD/DVD Image**.
2. From the **Keyboard** menu, click **Ctrl+Alt+Del** to start the system.
3. When “Press or <F2> to enter setup” displays during system startup, press F2.

The Extensible Firmware Interface (EFI) setup screen opens.

4. From the EFI menu, select **Boot**.
5. From the Boot Option Priorities area, select the boot device **HITACHI Remote CD/DVD** for **Boot Option #1**.
6. From the EFI menu, select **Save & Exit**, and then **Save Changes**.

Install Oracle Enterprise Linux 5, Update 5, on Oracle Database Server

This procedure assumes that the configuration of the Oracle database server is to boot the operating system from an ISO image on a remote CD/DVD image. See [Mount the Oracle Enterprise Linux 5, Update 5, ISO Image from a Remote CD/DVD Image](#).

The Oracle database server starts with a splash screen for Oracle Enterprise Linux. The bottom of the screen lists instructions, function keys, and the boot prompt.

To install Oracle Enterprise Linux 5U5 on Oracle database server, do the following.

1. Boot Hitachi Compute Blade 2000.
 2. Display the installation instructions.
 - On the Oracle Enterprise Linux splash screen, at the **boot** prompt, type the following:

```
l i n u x  t e x t
```
 3. Follow the on-screen instructions to complete the installation.
 4. Change to boot the Compute Blade 2000 server from the Hitachi Unified Storage VM storage system.
 - (1) Reboot the system.
 - (2) When “Press or <F2> to enter setup” displays during system startup, press F2.
 - (3) From the Boot Option Priorities area, select **Hitachi Unified Storage VM** for **Boot Option #1**.
 - (4) From the EFI menu, select **Save & Exit**, and then **Save Changes**.
-

Install and Configure Oracle Enterprise Linux 5, Update 5, Device Mapper

This solution uses the round-robin load-balancing algorithm. Your choice of load-balancing algorithm depends on the specific environment and access patterns of the application. In some environments, such as an environment shared with other applications, one of the other algorithms might give the best overall performance.

To install and configure Oracle Enterprise Linux 5, Update 5, Device Mapper, follow the instructions in *DM Multipath*. This document is in [Red Hat Customer Portal](#) on the **Knowledge** tab in HTML and PDF.

Install and Configure Fusion-io ioDrive2 PCI-e Flash Card Acceleration (Optional)

Note — Using Fusion-io ioDrive2 PCI-e flash card acceleration in your installation is optional.

When using flash acceleration, this solution uses four 1.2 TB Fusion-io PCI Express (PCI-e) flash cards.

Install Fusion-io ioDrive2 Storage Devices

To Install Fusion-io ioDrive2s in the Oracle database server, do the following.

1. Shut down the Oracle database server.
 2. Click **Launch Remote Console** and log on to Remote Console.
User Name: user01
Password: pass01
Remote Console opens.
 3. Power off the server blades.
 - From the **Power and Reset** menu, click **Forced Power Off**.
 4. Install the Fusion-io ioDrive2 storage devices in the PCI-e expansion Slot 0 and PCI-e expansion Slot 1 for each server blade.
 - This solution uses two server blades and four Fusion-io ioDrive2s.
 5. Power on the server blades.
 - From the **Power and Reset** menu, click **Power On**.
-

Configure Fusion-io ioDrive2 Storage Devices

You need an account on [Fusion-io Customer Support](#) to complete this process.

To configure Fusion-io ioDrive2 storage devices, do the following.

1. Download the driver for the Fusion-io ioDrive2 from Fusion-io Customer Support.
2. Follow the instructions in the knowledge base on Fusion-io Customer Support to configure the driver.
3. Check the status of the Fusion-io ioDrive2s by typing the following command:

```
chkconfig --list iomemory-vsl
fi o-status -a
```

4. Check the path of the Fusion-io ioDrive2s by typing the following command:

```
ls -l /dev/fct*
```

- In case the Fusion-io ioDrive2s are not attached, attach the drives by typing the following command:

```
fi o-attach /dev/[drive number]
```

- For example, for drive /dev/fct0, type the following:

```
fi o-attach /dev/fct0
```

- This includes attaching the drives after the reboot of the Oracle database server

- Repeat this command for each Fusion-io ioDrive2:

- /dev/fct0

- /dev/fct1

- /dev/fct2

- /dev/fct3

After configuration, an ioDrive2 re-labels as a block device:

- /dev/fioa

- /dev/fiob

- /dev/fioc

- /dev/fiod

Install Oracle Database 11g Release 2 Automatic Storage Management

By installing Oracle Grid Infrastructure, you successfully create an instance of Oracle Database 11g Release 2 Automatic Storage Management (ASM).

To install Oracle Grid Infrastructure Software, follow the steps in [Oracle Grid Infrastructure Installation Guide](#).

Configure Oracle Database 11g Release 2 Automatic Storage Management

Use Without PCI-e Flash Acceleration or Using PCI-e Flash Acceleration depending on your installation.

Without PCI-e Flash Acceleration

Table 11 defines Oracle ASM disks and disk groups for on-line redo logs, control file, archived redo logs and Oracle database application data when not using PCI-e flash acceleration.

Table 11. Oracle ASM Disks and Disk Groups for Online Redo Logs, Control File, Archived Redo Logs and Application Data without PCI-e Flash Acceleration

LUN	LUN Size (GB)	OS device /dev/...	ASM disk	ASMDG	Purpose	Storage Ports
1	1000	360060e8013275e005020275e0000001p1	RGDISK01	REDODG01	<ul style="list-style-type: none"> ■ Online REDO log group and control file 	<ul style="list-style-type: none"> ■ 1A ■ 2A ■ 1B ■ 2B
2	1000	360060e8013275e005020275e0000002p1	RGDISK11	REDODG11	<ul style="list-style-type: none"> ■ Online REDO log group and control file 	<ul style="list-style-type: none"> ■ 1A ■ 2A ■ 1B ■ 2B
3	1000	360060e8013275e005020275e0000003p1	ARDISK01	ARCHDG	<ul style="list-style-type: none"> ■ Archived Redo Log file 	<ul style="list-style-type: none"> ■ 1A ■ 2A ■ 1B ■ 2B

Table 11. Oracle ASM Disks and Disk Groups for Online Redo Logs, Control File, Archived Redo Logs and Application Data without PCI-e Flash Acceleration (Continued)

<i>LUN</i>	<i>LUN Size (GB)</i>	<i>OS device /dev/...</i>	<i>ASM disk</i>	<i>ASMDG</i>	<i>Purpose</i>	<i>Storage Ports</i>
4 – 32	200	360060e8013275e005020275e00000022p1 - 360060e8013275e005020275e00000003ep1	DADISK01 – DADISK29	DATADG	<ul style="list-style-type: none"> ■ Sys ■ Undo ■ Temp ■ Application Data 	<ul style="list-style-type: none"> ■ 1A ■ 2A ■ 1B ■ 2B

To configure Automatic Storage Management when not using PCI-e flash acceleration, do the following.

1. Create Automatic Storage Management disk groups.
 - (1) Log on as a grid user from the bin directory located in Oracle Grid Infrastructure home.
 - Verify that you enable X Terminal. Follow the documentation on how to enable X Terminal for the X server software that you are using.
 - (2) Configure the Oracle Grid Infrastructure home to be the following values:


```
/u01/app/gri d/product/11. 2. 0/gri d
$ /u01/app/gri d/product/11. 2. 0/gri d/bi n/asmca
```
2. Follow the on-screen instructions to create the ASM disk groups.

Using PCI-e Flash Acceleration

Install and configure Fusion-io ioDrive2 PCI-e flash acceleration before completing this process. See *Install and Configure Fusion-io ioDrive2 PCI-e Flash Card Acceleration (Optional)*.

Table 12 lists Oracle ASM disks and disk groups for on-line redo logs, control file, and archived redo logs using PCI-e flash acceleration.

Table 12. Oracle ASM Disks and Disk Groups for Online Redo Logs, Control File, Archived Redo Logs and Application Data with PCI-e Flash Acceleration

<i>LUN</i>	<i>LUN Size (GB)</i>	<i>OS device / dev/...</i>	<i>ASM disk</i>	<i>ASMDG</i>	<i>Purpose</i>	<i>Storage Ports</i>
1	1000	360060e8013275e005020275e00000001p1	RGDISK01	REDODG01	<ul style="list-style-type: none"> ■ Online REDO log group and control file 	<ul style="list-style-type: none"> ■ 1A ■ 2A ■ 1B ■ 2B
2	1000	360060e8013275e005020275e00000002p1	RGDISK11	REDODG11	<ul style="list-style-type: none"> ■ Online REDO log group and control file 	<ul style="list-style-type: none"> ■ 1A ■ 2A ■ 1B ■ 2B
3	1000	360060e8013275e005020275e00000003p1	ARDISK01 – ARDISK03	ARCHDGD	<ul style="list-style-type: none"> ■ Archived Redo log file 	<ul style="list-style-type: none"> ■ 1A ■ 2A ■ 1B ■ 2B
N/A	1200	fioa – fioid	FIODISK1 – FIODISK4	PFMDG	<ul style="list-style-type: none"> ■ ASM Preferred Mirror Read with 1st Failure Group – FGFIO 	<ul style="list-style-type: none"> ■ 1A ■ 2A ■ 1B ■ 2B
4 - 27	200	360060e8013275e0000022p1 - 360060e8013275e005020275e00000030p1	DADISK01 – DADISK24		<ul style="list-style-type: none"> ■ ASM Preferred Mirror Read with 2nd Failure Group – FGSA 	<ul style="list-style-type: none"> ■ 1A ■ 2A ■ 1B ■ 2B
28 - 32	200	360060e8013275e005020275e0000003ap1 - 360060e8013275e005020275e0000003ep1	DADISK25 – DADISK29	DATADG	<ul style="list-style-type: none"> ■ Sys ■ Undo ■ Temp 	<ul style="list-style-type: none"> ■ 1A ■ 2A ■ 1B ■ 2B

To configure Automatic Storage Management when using PCI-e Flash Acceleration, do the following.

1. To map the Fusion-io ioDrive2s to Oracle ASM, type the following command for each Fusion-io ioDrive2:

```
/usr/sbin/asmtool -C -I /dev/oracleasm -n ASMFI0A -s /dev/[drive name] -a force=yes
```

- Repeat this command for each Fusion-io ioDrive2:

- /dev/fioa
- /dev/fiob
- /dev/fioc
- /dev/fiod

For example, for drive /dev/fioa, type the following:

```
/usr/sbin/asmtool -C -I /dev/oracleasm -n ASMFI0A -s /dev/fioa -a force=yes
```

2. To verify the labeling of the Fusion-io ioDrive2s as Oracle ASM disks, type the following command:

```
oracleasm lsdisks | grep -i fio
```

3. Follow the instructions to create disk groups as described in Without PCI-e Flash Acceleration.
4. Configure the ASM_PREFERRED_READ_FAILURE_GROUPS initialization parameter in the Oracle ASM instance to specify a list of failure groups as preferred read disks by typing the following:

```
ASM_PREFERRED_READ_FAILURE_GROUPS=PFMDG.FGFI0
```

Install Oracle Database 11g Release 2

To install Oracle Database 11g Release 2, follow the instructions in *Database Quick Installation Guide for Linux x86-64*.

This document in HTML and PDF is in [Installing and Upgrading](#) of the Oracle Database Documentation Library.

Create the Oracle Database

To create the Oracle database, you can use customized scripts or the dbca tool that is packaged in the Oracle software library. This solution uses the dbca tool.

Make sure you have enabled X Terminal. Follow the documentation on enabling X Terminal for the X server software that you are using.

To create the Oracle database, do the following:

1. Launch the dbca tool.
 - (1) Log on as Oracle User from the bin directory located in Oracle Software home.
 - (2) Type the following to configure the Oracle software home:

```
/u01/app/oracle/product/11.2.0/dbhome  
$/u01/app/oracle/product/11.2.0/dbhome/bin/dbca
```
 2. Follow the instructions in the wizard to create the Oracle database.
-

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.



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

© Hitachi Data Systems Corporation 2013. All rights reserved. HITACHI is a trademark or registered trademark of Hitachi, Ltd. "Innovate with Information" is a trademark or registered trademark of Hitachi Data Systems Corporation. Microsoft, Windows, and Windows Server are trademarks or registered trademarks of Microsoft Corporation. All other trademarks, service marks, and company names are properties of their respective owners.

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