

Deploying Microsoft Exchange Server 2010 on the Hitachi Virtual Storage Platform with Hitachi Dynamic Tiering

Implementation Guide

By Jeff Chen

May 2011



Feedback

Hitachi Data Systems welcomes your feedback. Please share your thoughts by sending an email message to SolutionLab@hds.com. Be sure to include the title of this white paper in your email message.

Table of Contents

Tested Solution Components	3
Hardware Components	3
Software Components.....	6
Solution Implementation	7
Configure the Hitachi Compute Blade 2000.....	10
Configure the SAN	18
Configure the Hitachi Virtual Storage Platform	21
Set Up SAN OS Boot	26
Install Microsoft Windows Server 2008 R2 on a Logical Partition.....	34
Install Hitachi Dynamic Link Manager Software.....	34
Install Exchange on Logical Partitions	34

Deploying Microsoft Exchange Server 2010 on the Hitachi Virtual Storage Platform with Hitachi Dynamic Tiering

Implementation Guide

This implementation guide focuses on deploying Microsoft Exchange Server 2010 using a Hitachi Compute Blade 2000 with logical partitioning (LPAR) technology. It is based on the reference architecture in the white paper *Deploying Microsoft Exchange Server 2010 on the Hitachi Virtual Storage Platform with Hitachi Dynamic Tiering Reference Architecture Guide*.

The intended use is by IT administrators responsible for Microsoft Exchange and storage. It assumes some familiarity with Hitachi Storage Navigator software, Microsoft Windows Server 2008 R2, and Exchange Server 2010.

IT administrators need email solutions that provide data protection and simple management in environmentally friendly data centers. Using Microsoft® Exchange Server® 2010 with storage area network (SAN) storage like the Hitachi Virtual Storage Platform accomplishes those business-critical objectives. This implementation guide uses changes in Exchange Server 2010 and the advantages of the Virtual Storage Platform to lower administration costs, improve system efficiency, and enable virtualization.

Exchange Server 2010 can be deployed for many types and numbers of users in a wide variety of infrastructure topologies. By implementing Exchange Server 2010 with the Virtual Storage Platform storage systems, you can effectively scale an environment from a few thousand users to hundreds of thousands of users.

Tested Solution Components

These are the deployed hardware and software components for this solution.

Hardware Components

Table 1. Hardware Components

<i>Hardware</i>	<i>Description</i>	<i>Version</i>	<i>Quantity</i>
Hitachi Virtual Storage Platform	6 × 8Gb/sec Fibre Channel ports used 128GB cache memory 96 × 300GB, 10K RPM, SAS disks 44 × 2TB 7.2K RPM, SATA disks	70-01-62-00/00	1

<i>Hardware</i>	<i>Description</i>	<i>Version</i>	<i>Quantity</i>
Hitachi Compute Blade 2000 Chassis	10 x 8Gb/sec dual-port HBAs 2 x 1Gb/sec LAN Switch Module 2 x management modules 8 x cooling fan modules 4 x power supply modules	A0154-E-5234	1
Hitachi Compute Blade 2000 — X55A2 Blade	2 x 6-Core Intel Xeon X5670 2.93GHz 72GB memory	58.22	4
Hitachi Compute Blade 2000 — X57A1 Blade	2 x 8-Core Intel Xeon X7560 2.26GHz 256GB memory	58.22	1
Brocade 5300 Fibre Channel Switch	8Gb/sec Fibre Channel ports	FOS 6.4.0E	2

Hitachi Compute Blade 2000 Chassis Configuration

This solution uses one Hitachi Compute Blade 2000 chassis with:

- Four standard X55A2 blades
- One high performance X57A1 blade
- Two standard 1Gb/sec LAN switch modules
- Ten Hitachi dual port 8Gb/sec HBA cards

Each blade has:

- Two on-board NICs. Each NIC is connected to a LAN switch module.
- Two PCIe slots available. All PCIe slots have Hitachi dual port 8Gb/sec HBA cards installed.

Hitachi HBAs are required to support the logical partitioning (LPAR) feature. With Hitachi HBAs and logical partitioning enabled, there are eight virtual WWNs created for the logical partitions on each blade. Figure 1 shows the front and back view of the Hitachi Compute Blade 2000 used in this solution.




Figure 1

Table 2 gives the details about the server blades used in this solution.

Table 2. Hitachi Compute Blade 2000 Configuration

<i>Blade Server</i>	<i>Logical Partitions (LPAR)</i>	<i>Server Name</i>	<i>Role</i>	<i>Number of CPU Cores</i>	<i>Memory (GB)</i>
Blade 1	LPAR2	AD2	Active Directory and DNS	4	8
Blade 2	LPAR1	MBX1	Mailbox Server	6	64
	LPAR2	CAS-HT1	Client Access and Hub Transport Server	6	8
Blade 3	LPAR1	MBX2	Mailbox Server	6	64
	LPAR2	CAS-HT2	Client Access and Hub Transport Server	6	8
Blade 4	LPAR1	MBX3	Mailbox Server	6	64
	LPAR2	CAS-HT3	Client Access and Hub Transport Server	6	8
Blade 5	LPAR1	MBX4	Mailbox Server	6	64
	LPAR2	CAS-HT4	Client Access and Hub Transport Server	6	8

Software Components

Table 3 lists detailed information about the software components used in this solution.

Table 3. Software Components

<i>Hardware</i>	<i>Version</i>
Hitachi Storage Navigator	Dependent on microcode version
Hitachi Dynamic Tiering	Dependent on microcode version
Hitachi Dynamic Link Manager	6.5
Microsoft Windows Server	2008 R2 Enterprise
Microsoft Exchange Server	2010 SP1
Microsoft Exchange Jetstress 2010	14.01.0225.017
Microsoft Exchange Load Generator 2010	14.01.0180.003

Hitachi Dynamic Tiering Software

Hitachi Dynamic Tiering eliminates the time-consuming manual processes of data classification and movement between storage tiers. This optimizes tiered storage usage while improving performance. It is available only on the Hitachi Virtual Storage Platform.

Most data is accessed rarely after creation. It should not be stored on your most expensive storage tier, but moved to a lower, less expensive storage tier. Defining where and how long data should reside at any point in its life cycle can be complex and problematic.

Many organizations use data tiers to manage application performance, manually provisioning space from several storage technologies with different performance and cost characteristics. Data specialists look to past usage patterns to determine how to configure tiers manually. The storage infrastructure cannot respond effectively to dynamic application and data use in this scenario. When usage patterns change rapidly, manually tiered storage systems produce less than optimal results.

Using existing features of Hitachi Dynamic Provisioning software, Hitachi Dynamic Tiering automates the management of multiple storage tiers as a single entity. Hitachi Dynamic Tiering automatically uses virtual volumes with embedded smart tiering that monitor access and moves data at the 42MB page level. Hitachi Dynamic Provisioning software manages the data seamlessly by saving the data into pages and Hitachi Dynamic Tiering automatically moves infrequently referenced pages to lower cost tiers of storage. Moving pages instead of entire data sets or files reduces the time and storage space required to migrate data.

Hitachi Dynamic Tiering automatically moves infrequently referenced pages to lower cost storage tiers. When there is frequent access to a page on a lower tier, Hitachi Dynamic Tiering automatically moves the page to a higher tier. Using this approach, Hitachi Dynamic Tiering improves the availability and performance of your storage systems and the applications using that storage.

Previously, the creation of each dynamic provisioning pool had to use one RAID level and one disk type. Hitachi Dynamic Tiering on the Virtual Storage Platform allows a single pool to contain tiers made up of multiple types of RAID groups and any type of disk, supported by the Virtual Storage Platform managing the various tiers within a dynamic provisioning pool automatically. This eliminates most user management of storage tiers within a storage system to maintain peak performance under dynamic conditions without storage administrator intervention.

Hitachi Dynamic Tiering inherits thin provisioning and wide striping functionalities from the Hitachi Dynamic Provisioning software program. This provides virtual storage capacity to eliminate application service interruptions, reduce costs, and simplify administration:

- Optimize storage performance and capacity based on business or application requirements.
- Support deferring storage capacity upgrades to align with actual business usage.
- Simplify the storage administration process.
- Provide performance improvements through automatic optimized wide striping of data across all available disks in a storage pool.
- Eliminate hot spots across the different RAID groups by smoothing the combined workload.
- Improve capacity utilization significantly.

Use Hitachi Dynamic Tiering Software to ensure that the right data is in the right place at the right time. This maximizes service levels and minimizes total cost of storage ownership. For more information, see [Hitachi Dynamic Tiering Software](#).

Hitachi Dynamic Link Manager Software

The Hitachi Dynamic Link Manager software was used for SAN multipathing, configured with the round-robin multipathing policy. The round robin load balancing algorithm in Hitachi Dynamic Link Manager software automatically selects a path by rotating through all available paths, thus balancing the load across all available paths and optimizing IOPS and response time.

Solution Implementation

Before implementing this solution, make sure that the following items are available:

- USB DVD drive (comes with Hitachi Compute Blade 2000)
- Microsoft Windows 2008 installation DVD or ISO image
- Driver kit version 12-01 CD or ISO image (includes driver for Hitachi HBA)
- 20 IP addresses:
 - 1 for the management module
 - 5 for blades
 - 5 for logical partitioning control
 - 9 for guest operating systems on logical partitions
- 3 IP addresses in environments where these items are not already in place and connected:
 - 1 for the storage system
 - 2 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.

For high availability purposes, the storage area network (SAN) configuration for this implementation guide uses two fibre channel switches.

- Configure four redundant paths from the server blades to the storage ports for Exchange volumes
- Configure two redundant paths from the server blades to the storage ports for SAN OS boot volumes.

Figure 2 shows the fibre channel connections used for this solution.

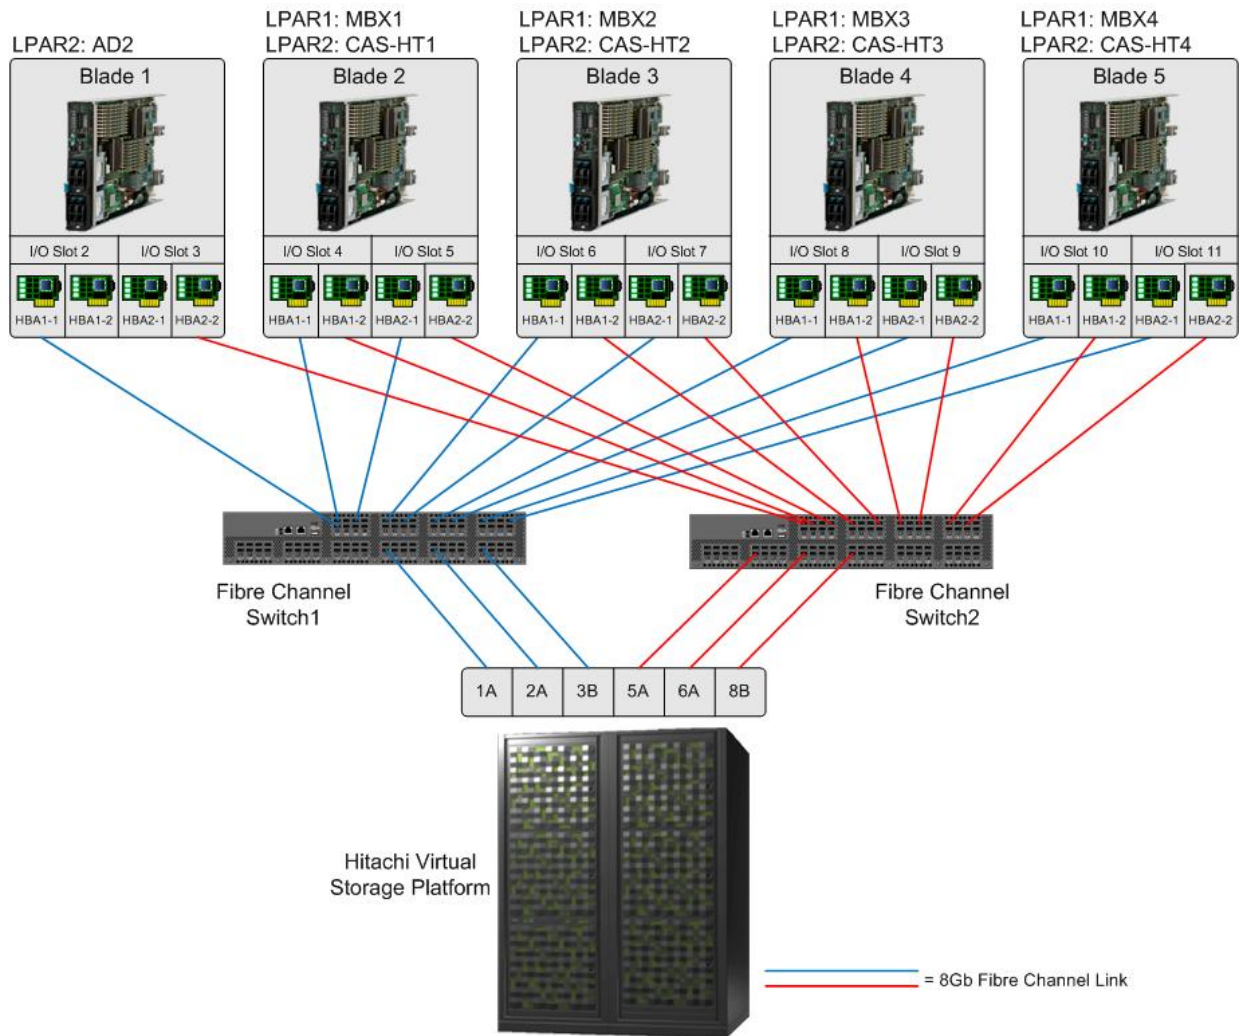


Figure 2

Deploy this Exchange Server 2010 solution following these high-level steps:

1. Configure the Hitachi Compute Blade 2000.
2. Configure the SAN.
3. Configure storage.
4. Set up SAN OS boot.
5. Install Windows 2008 on logical partitions.
6. Install Active Directory.
7. Install Hitachi Dynamic Link Manager software.
8. Install Exchange on logical partitions.

Some of these high-level tasks require that you use the following Hitachi Compute Blade 2000 interfaces:

- **Management module web GUI** — Manages entire chassis using a web-based GUI. To access it, open a browser and enter management module's IP address.
- **Blade server web GUI** — Manages a blade using a web-based GUI. To access it, open a browser and enter a blade's IP address.
- **Blade server command-line interface** — Manages the logical partitioning feature using a command-line interface. To access it, use a terminal emulator client to open a Telnet session to a blade's IP address.
- **Remote console** — Remote KVM for a blade. To access it, open a blade server web GUI. When the log on window displays, click **Launch Remote Console**.

Note — Hitachi is in the process of changing the product names for its Hitachi Compute Blade products. Some of these interfaces might use previous product names.

For more information about each of these tasks, see the following:

- Hitachi Storage Navigator online help
- Hitachi Dynamic Link Manager Software user's guide
- Hitachi Compute Blade 2000 user's guide
- Microsoft *TechNet* article "[Deploying Exchange 2010.](#)"

Configure the Hitachi Compute Blade 2000

This installation uses a Hitachi Compute Blade 2000 for Datacenter A and another for Datacenter B. Follow these procedures for each datacenter.

These procedures assume the following:

- The blade chassis is racked.
- The blades are installed into the chassis.
- All cabling for the LAN and SAN is complete.

Configure Required Management IP Addresses

Configure the IP addresses for the management module and the blades using the management module browser interface following these steps:

1. Connect a system console (laptop or personal computer) to the MGMT0 port of the management module using an Ethernet cable.

If two management modules are installed, connect to the management module with the MSR LED that is lit solid green.




Figure 3

2. Open a web browser, type this URL in the address bar and press ENTER:
http://192.168.0.1/

The log on screen displays on the browser.

3. Log on using the following:
 - User name (default): **admini strator**
 - Password (default): **password**

The management module interface opens.

4. Click the **Settings** tab.
5. From the left navigation bar, click **The configuration of network link**.

The **Management LAN network** pane opens.

6. Scroll down to click **Edit**.

The fields become editable.

7. Type an IP address, subnet mask, and default gateway in **The management module** area. Also, type an IP address, subnet mask and default gateway in the **Partition** area for each active partition (blade).

The management module and each active partition (blade) require its own IP address.

8. Click **Confirm**, and then click **Apply**.

This saves the IP addresses.

9. Disconnect the Ethernet cable from the laptop or personal computer to the management module.

10. Connect an Ethernet cable from the management module to an external management LAN switch

You use this cable to perform management through any system console in the management LAN in "Set Up HVM."

Configure Blades

Complete the following for each active blade.

Set Up HVM

This procedure requires the connection to the management LAN from step 10 in “Configure Required Management IP Addresses.”

To set up HVM, follow these steps:

1. Enter the blade IP address in the address bar of a browser.
2. Log on using the following:
 - User name (default): **user01**
 - Password (default): **pass01**The blade web interface opens.
3. Click the **EFI Setup** link.
The **EFI Setup** pane opens.
4. For **SMT (Simultaneous Multi-Threading)**, select **Disable**.
5. For **PCI Error Handling Mode**, select **Legacy**.
6. Keep the default values for the other EFI settings, and click **Modify**.
The **EFI Setup (Confirm)** pane opens.
7. Click **Confirm**.
This updates the EFI settings.
8. Click the **HVM Setup** link.
The **HVM Setup** window opens.
9. For **OS Mode**, select **HVM Mode**.
10. Click **Modify**.
A confirmation window displays.
11. Click **Confirm**.
The confirmation window closes and the **HVM Setup** window opens.
12. Click the **Power and LEDs** link.
The **Power and LEDs** window opens.
13. Click **Power on**.
The blade is powered on.

Configure HVM

Configure HVM following these steps:

1. Start a Telnet session to the IP address assigned to the partition.

2. Log on using the following:

- User name (default): **admini strator**
- Password (default): **password**

The **System Configuration** window opens with a message that indicates the initialization of HVM.

3. After the initialization message closes, place the cursor on **HVM IP Address** and press ENTER.

The **Change of HVM IP Address** menu displays.

4. Type the IP address and press ENTER.
5. Place the cursor on **Subnet Mask** and press ENTER.

The **Change of Subnet Mask** menu displays.

6. Type the subnet mask and press ENTER.
7. Place the cursor on **Default Gateway** and press ENTER.

The **Change of Default Gateway** menu displays.

8. Type the IP address for the default gateway and press ENTER.
9. Place the cursor on **VNIC System No** and press ENTER.

The **VNIC System No Setting** menu displays.

10. Type a unique number from 1 to 128 and press ENTER.

The system uses the VNIC system number to create a unique MAC address that does not conflict with other MAC addresses. Use a unique value from 1 and 128 for each HVM.

11. Press F10.

A confirmation menu opens.

12. Place the cursor on **Yes** and press ENTER.

A message indicating that the configuration is changing displays. In about 40 seconds, a **Normal End** window displays.

13. Press ENTER.

The window closes.

14. Press ESC .

The main menu displays. Keep the Telnet session open for "Configure a Logical Partition."

Configure a Logical Partition

Configuring a logical partition (LPAR) involves assigning resources to it, including CPU, memory, network, and fibre channel. This procedure uses the Telnet session launched in “Configure HVM.” Complete this procedure for *each* logical partition.

To set up a logical partition, follow all these steps:

Name the logical partition:

1. Place the cursor on **Logical Partition Configuration** and press ENTER.
The **Logical Partition (LPAR) Configuration** menu displays.
2. Press F6.
The **Add LPAR** menu displays.
3. Place the cursor on the first **NO_NAME** and press ENTER.
The first line of the **Logical Partition (LPAR) Configuration** pane is populated.
4. Place the cursor on **NO_NAME** and press ENTER.
The **Logical Partition Name** window displays.
5. Type a name for the logical partition and press ENTER.

Set number of CPUs and memory size:

1. Place the cursor on the **Shr** column for shared CPU mode, and press ENTER.
The **number of Shared Logical Processors** pane displays.
2. Type the number of CPUs for the logical partition and press ENTER.

- Place the cursor on the **Mem** column and press ENTER.

The memory size (in MB) menu displays.

```

+-----+
|+ Logical Partition (LPAR) Configuration -----|
|| # Name   Sta Pro Shr Ded Srv   Mem VN ID AA AC PC VC EB   || | |
|| 1 LPAR1  Dea  2  0  2 100   1024 0  Y *  N *  N BIOS  ||
|| 2 LPAR2  Dea  1  0 -----+ *  N BIOS  ||
|| 3 LPAR3  Dea  1  0| The memory size (in MB) | *  N BIOS  ||
|| 4 LPAR4  Dea  1  0|                               | *  N BIOS  ||
|| 5                               | +1024          |           ||
|| 6                               |           |           ||
|| 7                               | -256 002048 +256 |           ||
|| 8                               |           |           ||
|| 9                               | -1024          |           ||
|| 10                              |           |           ||
||                               | Maximum size   : 4608 |p / [PageDown]:Page Down ||
|-----+-----+ Effective max   : 4608 |-----+-----|
|+ Logical Information -----+-----+ Physical Information -----+
||                               Pro Shr| Within bounds |er Memory   : 4608 MB ||
|| Assign Total   5  0|-----+processors   : 8(8)  ||
|| Act Total      0  0| F1:Input number in GB | Shared     : 0      ||
|| Remain -----+ Dedicate   : 0      ||
|-----+-----+-----+
| The memory size in Mega Byte
+-----+
| F2:Mem Alloc Dep F3:Act F4:Deact F5:React F6:Add F7:Remove Esc:Menu |
+-----+

```

Figure 4

- Press F1.

The **Input the memory size in GB** prompt displays.

- Type the memory size in gigabytes and press ENTER.

The menu closes.

Set auto activation and PCI device assignment:

- Place the cursor on an asterisk (*) in the **AA** (auto activation) column and press ENTER.

The auto activation order menu displays. When setting auto activation, each logical partition powers on in the order you specify in this column when you power on the HVM.

- Type the activation order number and press ENTER.

The menu closes.

- Press ESC.

The main menu displays.

- Place the cursor on **PCI Device Assignment** and press ENTER.

The **PCI Device Assignment** pane displays.

- Place the cursor on existing entry in the **Schd** row for a PCI device number and press ENTER.

The **PCI Device Scheduling mode Assignment** menu displays. This is where you specify that you want to use the shared NIC and fibre channel ports.

```

+-----+
|+ PCI Device Assignment -----+
||
||      PCI Device#: 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
||      Type:      U N N N N N
||      Schd:      E S+ S+ S+ S+ S+
||
|| # Name      Sta
|| 1 LEAR1     Dea   A - - - - -
|| 2 LEAR2     Dea |-----+
|| 3 LEAR3     Dea | PCI Device Scheduling mode Assignment |
|| 4 LEAR4     Dea |-----+
|| 5           |
|| 6           |           D
|| 7           |           S
|| 8           |-----+
|| 9           | PCI Device is Shared Mode
|| 10          |-----+
||
||                                     [PageUp]:Page Up / [PageDown]:Page Down
+-----+
|+ Selected PCI Device Information -----+
|| # Vendor      Device Name          Slot# Bus# Dev# Func#
|| 2 Intel Corp.  GbE Controller             E00  62  0  0
+-----+
| F5:Attach/Detach  F10:Update PCI Dev Schd  F11:Left  F12:Right  Esc:Menu
+-----+

```

Figure 5

- Place the cursor on **S** for shared mode and press ENTER.

The menu closes.

- Repeat Step 5 and Step 6 for every PCI device.
- Press F10 to save the settings.

The **Save Setting** window displays.

- Place the cursor on **Yes** and press ENTER.

A message indicating that the configuration is being changed displays. After a few minutes, a **Normal End** message displays.

- Press ENTER.

The message closes.

- Press ESC.

The main menu displays.

Set VNIC assignment:

- Place the cursor on **VNIC Assignment** and press ENTER.

The **Virtual NIC Assignment** pane displays.

- Put cursor on the asterisk (*) below **Virtual NIC Number 0** and press ENTER.

The **Physical NIC/Port Number Setting** menu displays.

- Place the cursor on **1a** and press ENTER.

The Virtual NIC Number 0 and Virtual NIC Number 1 are set to 1a and 1b.

- Press ESC.

The main menu displays.

Assign the shared Fibre Channel (FC):

- Place the cursor on **Shared FC Assignment** and press ENTER.

The **Shared FC Assignment** menu displays.

- Place the cursor on an asterisk (*) and press ENTER.

The **Shared FC Assignment** menu displays.

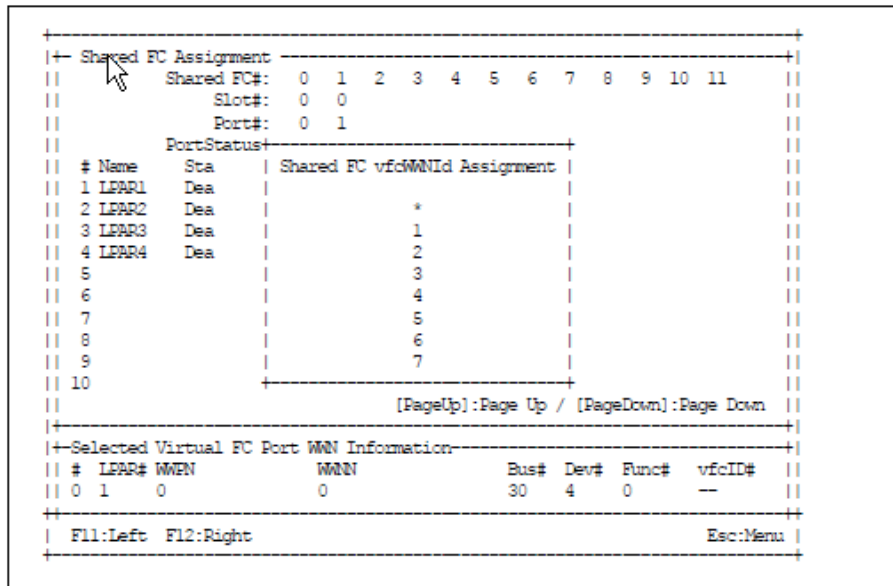


Figure 6

- Place the cursor on the first number available on the list and press ENTER.
- Repeat Step 27 and Step 28 for all shared fibre channel numbers.
- Press ESC.

The main menu displays.

- Press F9.

The message indicating that the configuration is being saved displays. It can display for several minutes. The configuration is saved when the **HVM configuration is saved successfully** message displays.

7. Press ENTER.

The message closes. Keep the Telnet session for “Collect WWN Information for Zoning.”

Collect WWN Information for Zoning

Enabling the HVM mode creates eight virtual WWNs for each Fibre Channel port.

This procedure uses the Telnet session from “Configure a Logical Partition.” To collect WWN information, follow these steps:

1. Place the cursor on **Allocated FC Information** and press ENTER.

The **Allocated FC Information** pane displays.

```

+- Allocated FC Information ----- vfc Seed Info. 1534
  Lpar# Slot# Port# schMd vfcID WWPN          WWNN
  1      4      0      S      1      2310000087bfc002 2310000087bfc003
  1      4      1      S      1      2310000087bfc022 2310000087bfc023
  1      5      0      S      1      2310000087bfc082 2310000087bfc083
  1      5      1      S      1      2310000087bfc0a2 2310000087bfc0a3
  2      4      0      S      2      2310000087bfc004 2310000087bfc005
  2      4      1      S      2      2310000087bfc024 2310000087bfc025
  2      5      0      S      2      2310000087bfc084 2310000087bfc085
  2      5      1      S      2      2310000087bfc0a4 2310000087bfc0a5
  --     4      0      S      3      2310000087bfc006 2310000087bfc007
  --     4      0      S      4      2310000087bfc008 2310000087bfc009
  --     4      0      S      5      2310000087bfc00a 2310000087bfc00b
  --     4      0      S      6      2310000087bfc00c 2310000087bfc00d
  --     4      0      S      7      2310000087bfc00e 2310000087bfc00f
  --     4      0      S      8      2310000087bfc010 2310000087bfc011
  --     4      0      S      9      2310000087bfc012 2310000087bfc013
  --     4      0      S     10     2310000087bfc014 2310000087bfc015
                                     [PageUp]:Page Up / [PageDown]:Page Down
-----
                                           Esc:Menu
  
```

2. Record the WWN information for later use.

You use this information when you zone SAN switches and create host groups on the storage system.

3. Press ESC.

The main menu displays. Leave the Telnet session open for “Activate a Logical Partition” in “Configure the SAN.”

Configure the SAN

Use ports 0A, 1A, 0E, and 1E on the Hitachi Virtual Storage Platform for Microsoft Exchange databases and logs.

Use ports 0C and 1C for SAN OS boot volumes.

Use Hitachi Dynamic Link Manager software for multipathing with the round-robin load balancing algorithm.

The following sections generally describe how to configure the SAN for this solution. These are specific configuration details for the Hitachi Data Systems lab environment. These details can vary for a production environment.

Zone SAN Switches

This Hitachi Data Systems lab used two Brocade 5300 SAN switches. If your environment uses different switches, the following steps can vary.

Configure two zones for each host (SAN OS boot). Each zone must use a different switch to provide redundancy. Configure four zones for each Microsoft Exchange mailbox host (Exchange database and log) to use different switches to provide redundancy.

To configure zoning for the SAN switches, follow these high-level steps:

1. Create aliases for ports.
2. Create zones.
3. Add members (aliases) to the zones.
4. Save the zones.
5. Change and save the zone configuration
6. Enable the zones.

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

Table 4 lists the zoning used for Datacenter A in the Hitachi Data Systems lab.

Table 4. Zoning Configuration

<i>Blade</i>	<i>Logical Partition Host</i>	<i>Virtual HBA Port</i>	<i>Storage System Port</i>	<i>Storage System Host Group</i>	<i>Zone Name</i>	
Blade 1	LPAR2 AD2	HBA1_1	3B	CB2K01B1L2H11	CB2K01_B1_L2_HBA1_1_VSP36_3B	
		HBA2_2	8B	CB2K01B1L2H12	CB2K01_B1_L2_HBA2_2_VSP36_8B	
Blade 2	LPAR1 MBX1	HBA1_1	1A	CB2K01B2L1H11	CB2K01_B2_L1_HBA1_1_VSP36_1A	
		HBA1_2	6A	CB2K01B2L1H12	CB2K01_B2_L1_HBA1_2_VSP36_6A	
		HBA2_1	2A	CB2K01B2L1H21	CB2K01_B2_L1_HBA2_1_VSP36_2A	
		HBA2_2	5A	CB2K01B2L1H22	CB2K01_B2_L1_HBA2_2_VSP36_5A	
		HBA1_1	3B	CB2K01B2L1H11	CB2K01_B2_L1_HBA1_1_VSP36_3B	
		HBA2_2	8B	CB2K01B2L1H22	CB2K01_B2_L1_HBA2_2_VSP36_8B	
	LPAR2 CAS-HT1	HBA1_1	1A	CB2K01B2L2H11	CB2K01_B2_L2_HBA1_1_VSP36_1A	
		HBA1_2	6A	CB2K01B2L2H12	CB2K01_B2_L2_HBA1_2_VSP36_6A	
		HBA2_1	2A	CB2K01B2L2H21	CB2K01_B2_L2_HBA2_1_VSP36_2A	
		HBA2_2	5A	CB2K01B2L2H22	CB2K01_B2_L2_HBA2_2_VSP36_5A	
		HBA1_1	3B	CB2K01B2L2H11	CB2K01_B2_L2_HBA1_1_VSP36_3B	
		HBA2_2	8B	CB2K01B2L2H22	CB2K01_B2_L2_HBA2_2_VSP36_8B	
	Blade 3	LPAR1 MBX2	HBA1_1	1A	CB2K01B3L1H11	CB2K01_B3_L1_HBA1_1_VSP36_1A
			HBA1_2	6A	CB2K01B3L1H12	CB2K01_B3_L1_HBA1_2_VSP36_6A
HBA2_1			2A	CB2K01B3L1H21	CB2K01_B3_L1_HBA2_1_VSP36_2A	

<i>Blade</i>	<i>Logical Partition Host</i>	<i>Virtual HBA Port</i>	<i>Storage System Port</i>	<i>Storage System Host Group</i>	<i>Zone Name</i>
		HBA2_2	5A	CB2K01B3L1H22	CB2K01_B3_L1_HBA2_2_VSP36_5A
		HBA1_1	3B	CB2K01B3L1H11	CB2K01_B3_L1_HBA1_1_VSP36_3B
		HBA2_2	8B	CB2K01B3L1H22	CB2K01_B3_L1_HBA2_2_VSP36_8B
	LPAR2 CAS-HT2	HBA1_1	1A	CB2K01B3L2H11	CB2K01_B3_L2_HBA1_1_VSP36_1A
		HBA1_2	6A	CB2K01B3L2H12	CB2K01_B3_L2_HBA1_2_VSP36_6A
		HBA2_1	2A	CB2K01B3L2H21	CB2K01_B3_L2_HBA2_1_VSP36_2A
		HBA2_2	5A	CB2K01B3L2H22	CB2K01_B3_L2_HBA2_2_VSP36_5A
		HBA1_1	3B	CB2K01B3L2H11	CB2K01_B3_L2_HBA1_1_VSP36_3B
		HBA2_2	8B	CB2K01B3L2H22	CB2K01_B3_L2_HBA2_2_VSP36_8B
Blade 4	LPAR1 MBX3	HBA1_1	1A	CB2K01B4L1H11	CB2K01_B4_L1_HBA1_1_VSP36_1A
		HBA1_2	6A	CB2K01B4L1H12	CB2K01_B4_L1_HBA1_2_VSP36_6A
		HBA2_1	2A	CB2K01B4L1H21	CB2K01_B4_L1_HBA2_1_VSP36_2A
		HBA2_2	5A	CB2K01B4L1H22	CB2K01_B4_L1_HBA2_2_VSP36_5A
		HBA1_1	3B	CB2K01B4L1H11	CB2K01_B4_L1_HBA1_1_VSP36_3B
		HBA2_2	8B	CB2K01B4L1H22	CB2K01_B4_L1_HBA2_2_VSP36_8B
	LPAR2 CAS-HT3	HBA1_1	1A	CB2K01B4L2H11	CB2K01_B4_L2_HBA1_1_VSP36_1A
		HBA1_2	6A	CB2K01B4L2H12	CB2K01_B4_L2_HBA1_2_VSP36_6A
		HBA2_1	2A	CB2K01B4L2H21	CB2K01_B4_L2_HBA2_1_VSP36_2A
		HBA2_2	5A	CB2K01B4L2H22	CB2K01_B4_L2_HBA2_2_VSP36_5A
		HBA1_1	3B	CB2K01B4L2H11	CB2K01_B4_L2_HBA1_1_VSP36_3B
		HBA2_2	8B	CB2K01B4L2H22	CB2K01_B4_L2_HBA2_2_VSP36_8B
Blade 5	LPAR1 MBX4	HBA1_1	1A	CB2K01B5L1H11	CB2K01_B5_L1_HBA1_1_VSP36_1A
		HBA1_2	6A	CB2K01B5L1H12	CB2K01_B5_L1_HBA1_2_VSP36_6A
		HBA2_1	2A	CB2K01B5L1H21	CB2K01_B5_L1_HBA2_1_VSP36_2A
		HBA2_2	5A	CB2K01B5L1H22	CB2K01_B5_L1_HBA2_2_VSP36_5A
		HBA1_1	3B	CB2K01B5L1H11	CB2K01_B5_L1_HBA1_1_VSP36_3B
		HBA2_2	8B	CB2K01B5L1H22	CB2K01_B5_L1_HBA2_2_VSP36_8B
	LPAR2 CAS-HT4	HBA1_1	1A	CB2K01B5L2H11	CB2K01_B5_L2_HBA1_1_VSP36_1A
		HBA1_2	6A	CB2K01B5L2H12	CB2K01_B5_L2_HBA1_2_VSP36_6A
		HBA2_1	2A	CB2K01B5L2H21	CB2K01_B5_L2_HBA2_1_VSP36_2A
		HBA2_2	5A	CB2K01B5L2H22	CB2K01_B5_L2_HBA2_2_VSP36_5A
		HBA1_1	3B	CB2K01B5L2H11	CB2K01_B5_L2_HBA1_1_VSP36_3B
		HBA2_2	8B	CB2K01B5L2H22	CB2K01_B5_L2_HBA2_2_VSP36_8B

Configure the Hitachi Virtual Storage Platform

This solution uses two dynamic provisioning pools for databases and logs. The Multi-Tier Pool option (Hitachi Dynamic Tiering) is enabled on the database pool to accommodate multiple SAS RAID-5 (7D+1P) groups and SATA RAID-5 (7D+1P) groups within a single pool.

Create a Dynamic Provisioning Pool

To create a dynamic provisioning pool using Hitachi Storage Navigator software, follow these steps:

1. On the **Actions** menu, point to **Pool**, and then click **Create Pools**.

The **Create Pools** dialog box opens.

2. In the **Pool Type** list, click **Dynamic Provisioning**.
3. For **Multi-Tier Pool**, select **Disable**.
4. (Optional) Filter the available pool volumes by clicking a value in the **Drive Type/RPM** list and the **RAID Level** list.
5. Click **Select Pool VOLS**.

The **Select Pool VOLS** dialog box opens.

6. Select one or more pool volumes from **Available Pool Volumes** and click **OK**.

Each selected pool volume displays in the **Total Selected Pool Volume** area of the **Create Pools** dialog box with a value in **Total Selected Capacity**.

7. Type a pool name prefix in **Prefix**.
8. (Optional) Choose an initial number for the pool name in **Initial Number**.
9. Click to expand **Options**.
10. Choose a pool id in **Pool ID**.
11. Choose a subscription limit in **Subscription Limit**.

This is the percentage of oversubscription for this pool that you allow in your environment.

12. Choose a value for the **Warning Threshold**.

The **Warning Threshold** value determines when to trigger a pool capacity warning alert..

13. Choose a value for the **Depletion Threshold**.

The **Depletion Threshold** value determines when to trigger a pool capacity depletion alert.

14. Click **Finish**.

The **Create Pools** dialog box displays.

15. Review the information for accuracy and click **Apply**.

Create a Dynamic Tiering Pool

To create a dynamic tiering pool, follow the previous steps in “Create a Dynamic Provisioning Pool,” except for Step 3. Instead, click **Enable** for the **Multi-Tier Pool** option. You will also have to select manual or automatic tier management. For this reference architecture, you need to select **Auto** tier management with a 4 hour monitoring cycle (at a minimum).

To add different drive type into a pool, follow these steps:

1. On the left **Explorer** panel, click **Storage Systems**.

The tree for the VSP displays.

2. From the VSP tree, click **Pools**, and then click the dynamic tiering pool.

The pool information displays on the right.

3. From the **Actions** menu, point to **Pool**, and then click **Expand Pool**.

The **Expand Pool** dialog box displays.

4. (Optional) Filter the available pool volumes by clicking a value in the **Drive Type/RPM** list and the **RAID Level** list.

5. Click **Select Pool VOLs**.

The **Select Pool VOLs** dialog box opens.

6. Select one or more pool volumes from **Available Pool Volumes** and click **Add**.

The selected volumes display in the **Selected Pool Volumes** area.

7. Verify the selected pool volumes and click **OK**.

The **Expand Pools** dialog box opens with values in **Total Selected Pool Volume** and **Total Selected Capacity**.

8. Click **Finish**.

The **Expand Pools** dialog box opens.

9. Click **Apply**.

Create a Host Group

Before creating a host group, enable port security. Port security isolates traffic from multiple servers and groups them in common configurations. Host group members must run the same operating system and connect to the same port.

You can define up to 2,048 LUN paths for one host group. You can assign up to 255 host groups and up to 255 WWPNs to one fibre channel port.

To create a host group using Storage Navigator software, follow these steps:

1. On the **Actions** menu, point to **Ports/Host Groups**, and then click **Create Host Groups**.

The **Create Host Groups** window opens.

2. Type a name in **Host Group Name**.
3. For **Host Mode**, select **2C[Windows]**.
4. In the **Available Hosts** area, select one or more hosts.
5. In the **Available Ports** area, select one or more ports.
6. Click **Add**.

There are values added to **Selected Host Groups**.

7. Click **Finish**.
The **Create Host Groups** window opens.
8. Click **Apply**.

Create a LDEV Within a Pool

To create a LDEV using Storage Navigator software, follow these steps:

1. On the **Actions** menu, point to **Logical Device**, and then click **Create LDEVs**.
The **Create LDEVs** window opens.
2. For **Provisioning Type**, select **Dynamic Provisioning**.
3. For **Emulation Type**, select **OPEN-V**.
4. (Optional) Filter the available pools by clicking a value in the **Drive Type/RPM** list and the **RAID Level** list.
5. Click **Select Pool**.

The **Select Pool** window opens.

6. Click a pool in **Available Pools** and click **OK**.
The **Create LDEVs** dialog box opens with values in **Selected Pool Name** and **Selected Pool Capacity**.
7. Enter a capacity amount in **LDEV Capacity** and click a unit of measure from the list.
8. Enter the number of LDEVs of that size to be created in **Number of LDEVs**.
9. In **LDEV Name**, choose a prefix in **Prefix** and choose an initial number in **Initial Number**.
10. Click to expand **Options**.
11. Review the value in **LDKC**. Modify the LDKC value if the default of 00 is not appropriate.
Most often, this is the case if the storage configuration has more than one LDKC.
12. Click a value from the **CU** list.
13. Click a value from the **DEV** list.

14. (Optional) Select a value from the **Interval** list. Leave the default value of **0** for sequential numbering of LDEVs.

If you want a different numbering sequence, choose a different value.

15. Review the default values in **Initial SSID**, **CLPR**, and **Processor Blade**.

In most situations, use the default values. Change them only if your environment requires different values.

16. Click **Add**.

Values are added to **Selected LDEVs**.

17. Click **Finish**.

The **Create LDEVs** window opens.

18. Click **Apply**.

Map an LDEV to a Host Group

To map an LDEV using Hitachi Storage Navigator software, follow these steps:

1. From the **Actions** menu, point to **Logical Device**, and then click **Add LUN Paths**.

The **Add LUN Paths** window opens.

2. In the **Available LDEVs** area, select one or more LDEVs.

3. Click **Add**.

Each selected LDEV displays in **Selected LDEVs**.

4. Click **Next**.

The **Add LUN Paths** window opens.

5. In the **Available Host Groups** area, select one or more host groups.

6. Click **Add**.

Each selected host group displays in **Selected Host Groups**.

7. Click **Next**.

The **Add LUN Paths** window opens.

8. Click **Finish**.

The **Add LUN Paths** window opens.

9. Click **Apply**.

Configure Storage for SAN OS Boot and Exchange

To satisfy 20,000 users needing 1GB of mailbox capacity and an I/O profile of 0.12 IOPS, this implementation guide uses one dynamic tiering pool and two dynamic provisioning pools. Pool 10 is configured as a dynamic tiering pool created from 12 SAS RAID-5 (7D + 1P) groups and 5 SATA RAID-5 (7D+ 1P) groups.

- Pool 10 contains 12 active databases and 12 passive databases.
- Pool 11 contains 12 active logs and 12 passive logs.
- Pool 00 contains 12 LUNs for logical partition OS boot volumes.

Table 5 shows configurations for each of the dynamic tiering and dynamic provisioning pools used in the Hitachi Data Systems lab.

Table 5. Pool Configuration

<i>Dynamic Provisioning Pool</i>	<i>Multi-Tier Pool (Dynamic Tiering)</i>	<i>Drive Type</i>	<i>RAID Configuration</i>	<i>Tier</i>	<i>Number of Drives</i>	<i>Drive Capacity (GB)</i>	<i>Usable Pool Capacity (TB)</i>
Pool 10	Enabled	SAS	RAID-5 (7D+1P)	1	96	300	22.01
		SATA	RAID-5 (7D+1P)	2	40	2,000	62.67
Pool 11	Disabled	SAS	RAID-5 (7D+1P)	N/A	24	300	5.50
Pool 00	Disabled	SATA	RAID-5 (3D+1P)	N/A	4	2,000	5.37

Table 6 shows the LUN allocation for each of the dynamic tiering and dynamic provisioning pools used in the Hitachi Data Systems lab.

Table 6. LUN Allocation

<i>LUN Allocation</i>	<i>Pool</i>	<i>LDEV Size (GB)</i>	<i>LDEVs</i>	<i>Storage Ports</i>
Database LUNs	Pool 10	2,048	10:01,10:02,10:03,10:04,10:05,10:06,10:07,10:08,10:09,10:0A,10:0B,10:0C 12:01,12:02,12:03,12:04,12:05,12:06,12:07,12:08,12:09,12:0A,12:0B,12:0C	CL1A,CL2A,CL5A,CL6A
Log LUNs	Pool 11	200	11:01,11:02,11:03,11:04,11:05,11:06,11:07,11:08,11:09,11:0A,11:0B,11:0C 13:01,13:02,13:03,13:04,13:05,13:06,13:07,13:08,13:09,13:0A,13:0B,13:0C	CL1A,CL2A,CL5A,CL6A
LPAR OS boot LUNs	Pool 00	127	05:06,05:07,05:08,05:09,05:10,05:11,05:12,05:13	CL3B,CL8B

Set Up SAN OS Boot

SAN boot is the capability to boot the OS from a LUN on a storage subsystem without using any of the server's internal disks. Utilizing logical partitions and the N+1 cold standby feature requires SAN boot. This describes how to set up the SAN OS boot.

Repeat these procedures for each logical partition.

Assign USB Devices and the Remote Console to a Logical Partition

Skip this procedure when setting up the first logical partition. By default, the USB devices and remote console are assigned to the first logical partition. Logical partitions can share USB ports and the remote console. However, they can be assigned to only one logical partition at a time. Before you set up a second logical partition on a blade, first assign the USB devices and the remote console to the second logical partition.

To assign USB devices and the remote console to a logical partition, follow these steps:

1. Start the blade server command-line interface.
2. On the main menu, place the cursor on **PCI Device Assignment** and press ENTER.

The **PCI Device Assignment** window displays.

3. Press F5.

The **Device Attach/Detach** menu displays.

4. Place the cursor on the target logical partition that you want to attach and press ENTER.

A confirmation message displays.

5. Place the cursor on **Yes** and press ENTER.

A message indicating that the change is being made displays. When finished, the PCI device is attached to the selected logical partition.

Activate a Logical Partition

This procedure assumes that the Telnet session from the "Collect WWN Information for Zoning" procedure is still open.

To activate a logical partition using the blade management command line interface, follow these steps:

1. On the main menu, press F3.

The **Activate LPAR** menu displays.

2. Place the cursor on an LPAR to be activated and press ENTER.

3. Place the cursor on **Continue (Don't show this message)** and press ENTER.

The selected logical partition is powered on. Leave the Telnet session open for "Configure HBAs."

Configure HBAs

This procedure assumes that the Telnet session from “Activate a Logical Partition” is still open.

To configure HBAs using the remote console, follow these steps:

Choose the port to configure:

1. Open a browser and enter the blade IP address in the address bar.

A log on screen displays.

2. Click **Launch Remote Console**.

The remote console log on window appears.

3. Log on using the following:

- User name (default): **user01**
- Password (default): **pass01**

The remote console launches.

4. Place the cursor on **Continue** and press ENTER.

The EFI shell launches.

5. Run this command: **drivers**

A list of drivers displays.

6. Look for **Hitachi PCI-X/PCIE Fibre channel DR** in the list of driver names. Write down the driver number in the left column labelled **DRV** for later reference.

In the example in Figure 7, the driver number is 70.

```
shell> drivers
          T  D
D          Y C I
R          P F A
V  VERSION  E G G #D #C DRIVER NAME                IMAGE NAME
== ===== = = = == == =====
36 0000000A D - - 3 - Platform Console Management Driver  ConPlatform
37 0000000A D - - 2 - Platform Console Management Driver  ConPlatform
.
.
.
70 10000118 D X - 4 - Hitachi PCI-X/PCIE Fibre channel Dr  ScsiBusFive
.
```

Figure 7

7. Run this command with the driver number from Step 6 at the command prompt:

```
drvcfg [driver number]
```

Using the example from Figure 7, the command looks like this:

```
drvcfg 70
```

A list of configurable components displays.

```
she]]> drvcfg 70
Configurable Components
Drv[70] Ctr1[87] Lang[eng]
Drv[70] Ctr1[88] Lang[eng]
Drv[70] Ctr1[8A] Lang[eng]
Drv[70] Ctr1[8B] Lang[eng]
```

Figure 8

8. Write down the driver number and controller number of the first configurable component on the list.
9. Run this command at the command prompt with the driver number from Step 6 and the controller number from Step 8:

```
drvcfg -s [driver number] [controller number]
```

Using the example from Figure 7 and Figure 8, the command looks like this:

```
drvcfg -s 70 7B
```

The command line prompt changes to this:

```
hfccfg>
```

10. Run this command at the command prompt:

```
select
```

A list of Fibre Channel ports and WWNs displays.

11. Type the port number that you plan to use for OS boot and press ENTER.

The command-line prompt changes to this:

```
hfccfg. [portwwn]>
```

This solution configured fibre channel Port 3.

12. Run this command at the command prompt:

```
set
```

A configuration dialog opens.

13. Verify that your settings match the following:

- **Boot function** — Enabled
- **Connection type** — Auto detection
- **Data rate** — Auto detection
- **Spinup delay** — Disabled
- **Login delay time** — 3 seconds
- **Persistent bindings** — Enabled
- **Force default parameter for adapter driver** — Disabled
- **Select boot device** — Enabled

Correct any setting that does not match the above.

Configure the first WWN in the boot device list:

1. At the **Change (y/n)** prompt, type: y

This prompt displays: Please select a number

2. To select a WWN to configure, type: 1

A list of options displays.

3. To scan for targets, type: 1

A list of target devices displays.

4. To scan for LUNs, type: 1

A list of LUNs displays.

5. To choose a LUN, type: 1

A list showing the target device's WWN and LUN displays. This prompt displays:
Update list#1 (y/n)

6. At the **prompt**, Type: y

The updated boot device list displays and a message indicating that the base settings are complete displays. You are returned to this prompt: hfcfg. [wwn] >

7. Run this command at the prompt: save

A confirmation prompt displays.

8. Type this to confirm: `y`

A confirmation message displays. The path to the storage for the boot OS is configured. The **hfccfg<wwn>>** prompt displays.

9. Run this command: `exit`

You are returned to the main EFI shell. Leave the remote session open for “Configure a Logical Partition to Boot Windows from an ISO Image on a Remote DVD.”

Map OS Boot Device

To map the OS boot device using EFI shell, follow these steps:

1. Run this command: `reconnect -r`

A success message displays.

2. Run this command: `map -r`

A device mapping table displays.

3. Run this command: `exit`

You are returned to the EFI main menu.

4. Place the cursor on **Boot Maintenance Manager** and press ENTER.

The Boot Maintenance Manager window displays.

5. Place the cursor on **Boot Options** and press ENTER.

The **Boot Maintenance Manager Boot Options** window displays.

6. Place the cursor on **Add Boot Option** and press ENTER.

The **File Explorer** window displays, as shown in Figure 9.

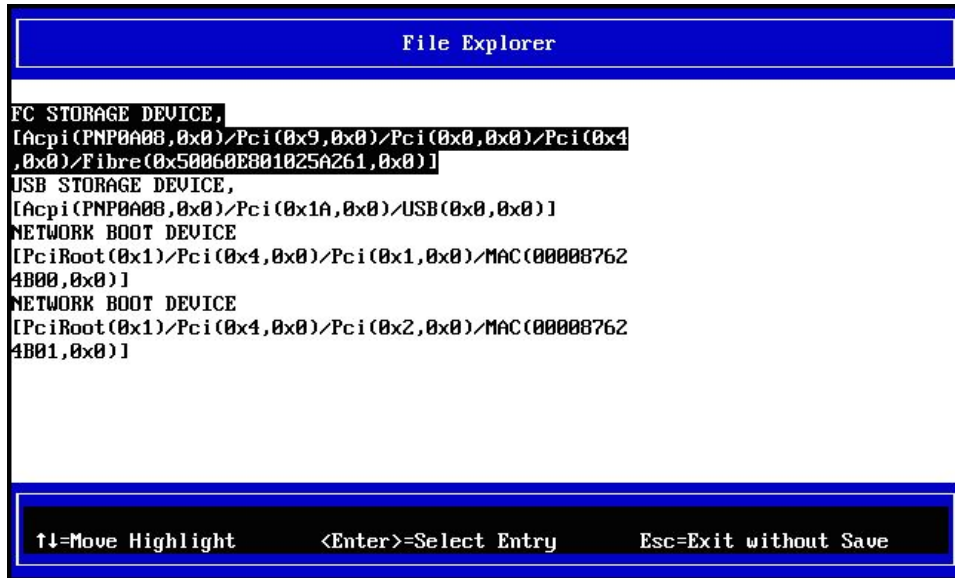


Figure 9

7. Place the cursor on **FC STORAGE DEVICE** and press ENTER.

Multiple input fields display.

8. Place the cursor on the **Input the description** field and press ENTER.

This field becomes editable.

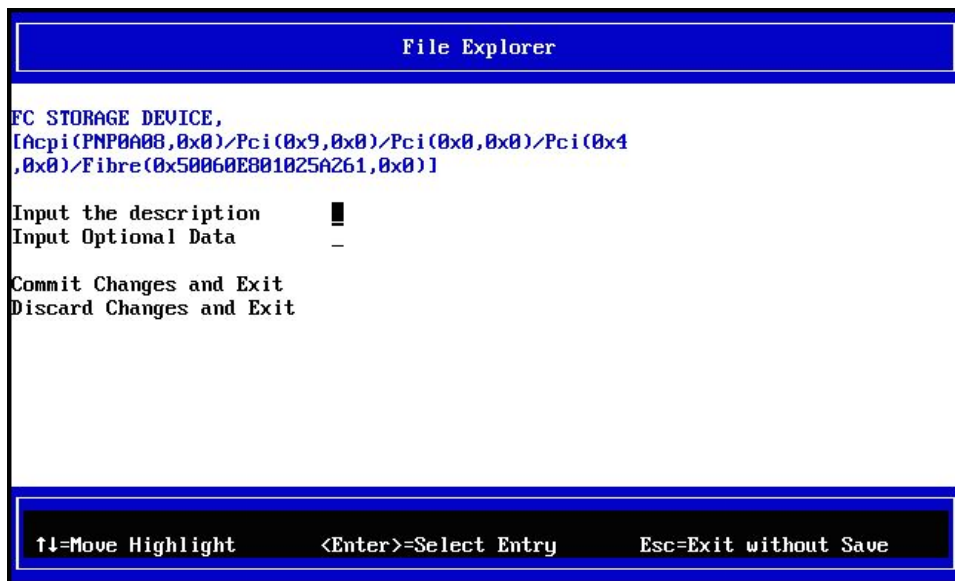


Figure 10

9. Type a name for the device in **Input the description** and press ENTER.

The prompt closes.

10. Place the cursor on **Commit Changes and Exit** and press ENTER.

You are returned to the EFI main menu.

Configure a Logical Partition to Boot Windows from an ISO Image on a Remote DVD

This procedure assumes that the remote console session from “Configure HBAs” is still open.

To mount the Windows 2008 R2 ISO image from a remote DVD, follow these steps:

1. From the **Remote CD/DVD** menu, click **Redirect CD/DVD Image**.

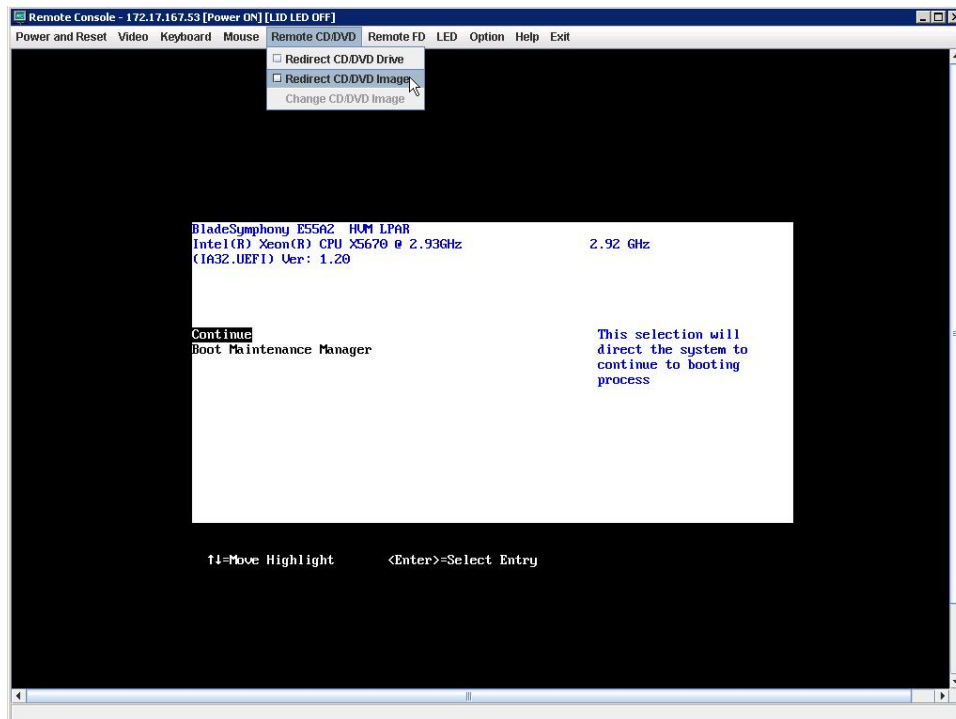


Figure 11

The **Open** window displays.

2. Navigate to your ISO image and click **Open**.

The **Open** window closes and you are returned to the EFI main menu.

3. Place the cursor on **Boot Maintenance Manager** and press ENTER.

The **Boot Maintenance Manager** window displays.

4. Place the cursor on **Boot Options** and press ENTER.

The **Boot Maintenance Manager Boot Options** window displays.

5. Place the cursor on **Add Boot Option** and press ENTER.

The **File Explorer** window displays.

6. Verify the listing of the USB storage device in the **File Explorer** window.

If the USB storage device is not listed, reactivate the logical partition and try again.

To reactivate the logical partition, go to the blade server command-line interface. On the main menu, press F5 to choose a logical partition to be reactivated. Soon after the reactivation, the EFI main screen displays on remote console. Return to Step 3 and repeat.

7. Place the cursor on **USB STORAGE DEVICE** and press ENTER.

Multiple input fields display.

8. Place the cursor on **Input the description** and press ENTER.

The field becomes editable.

9. Type a name for the device and press ENTER.

The prompt closes.

10. Place the cursor on **Commit Changes and Exit** and press ENTER.

You are returned to the **Boot Maintenance Manager** window.

11. Place the cursor on **Change Boot Order** and press ENTER.

The fields in the **Boot Maintenance Manager** window change.

12. Place the cursor on **Change the order** and press ENTER.

The **Boot Order** menu displays.

13. Use the plus (+) and minus (–) keys to change the boot order. Use the remote DVD first and the Fibre Channel storage LUN second. Then, press ENTER.

14. Place the cursor on **Commit Change and Exit** and press ENTER.

Your changes are saved and you returned to the **Boot Maintenance Manager** window.

15. Press ESC.

You are returned to the main menu.

16. Place the cursor on **Continue** and press ENTER.

The logical partition begins to boot from the Windows ISO image.

Install Microsoft Windows Server 2008 R2 on a Logical Partition

After the logical partition boots from the ISO image, follow the instructions from Microsoft to complete the installation. Specify the following:

- Hitachi X55X2 HBA driver
- Microsoft Windows Server 2008 R2 Enterprise Edition (full installation)
- A strong password for the local administrator account.

Install Windows Server 2008 R2 on all active logical partitions. Verify the installation of all Windows Server patches.

For more information, see the Microsoft TechNet article "[Installing Windows Server 2008 R2.](#)"

Install Hitachi Dynamic Link Manager Software

Hitachi Dynamic Link Manager software provides comprehensive multipathing, robust path failover, load balancing and integrated path management. To install Hitachi Dynamic Link Manager software, follow the instructions in the user guide that accompanies the software.

Install Exchange on Logical Partitions

Install the Exchange mailbox role on logical partition 1 on blade 2 through blade 5. Install client access and hub transport Exchange roles on logical partition 2 on blade 2 through blade 5.

For more information, see the Microsoft TechNet article "[Deploying Exchange 2010.](#)"

Create the DAG

You must create the DAG after installing and configuring the Microsoft Exchange mailbox servers.

To create the DAG for this environment, follow these high-level steps:

1. Create a DAG.
2. Add mailbox servers to the DAG.
3. Configure the DAG network.
4. Replicate active databases.

For more information, see the Microsoft TechNet article "[Managing Database Availability Groups.](#)"

Table 7 shows the DAG database distribution used in the Hitachi Data Systems lab.

Table 7. DAG Database Distribution

<i>Database</i>	<i>MBX1</i>	<i>MBX2</i>	<i>MBX3</i>	<i>MBX4</i>
DB1	Active	Passive		
DB2		Active	Passive	
DB3			Active	Passive
DB4	Passive			Active

<i>Database</i>	<i>MBX1</i>	<i>MBX2</i>	<i>MBX3</i>	<i>MBX4</i>
DB5	Active		Passive	
DB6		Active		Passive
DB7	Passive		Active	
DB8		Passive		Active
DB9	Active			Passive
DB10	Passive	Active		
DB11		Passive	Active	
DB12			Passive	Active

For more information about DAG database layout, see the Microsoft TechNet article "[Database Copy Layout Design.](#)"

 **Hitachi Data Systems Corporation**

Hitachi is a registered trademark of Hitachi, Ltd., in the United States and other countries. Hitachi Data Systems is a registered trademark and service mark of Hitachi, Ltd., in the United States and other countries. All other trademarks, service marks and company names mentioned in this document 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

© Hitachi Data Systems Corporation 2011. All Rights Reserved. AS-088-00 May 2011

Corporate Headquarters

750 Central Expressway,
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