

Deploying an 8,000-user Microsoft Exchange Server 2010 Environment with Hitachi Compute Blade 2000 on the Hitachi Adaptable Modular Storage 2300

Implementation Guide

By Jeff Chen and Leo Nguyen

February 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.....	4
Hardware Components	4
Software Components.....	6
Solution Implementation.....	7
Configure the Hitachi Compute Blade 2000.....	10
Configure the SAN	18
Configure Storage	20
Set Up SAN OS Boot	28
Install Windows 2008 R2 on an LPAR	35
Install Active Directory.....	35
Install Hitachi Dynamic Link Manager Software.....	35
Install Exchange on LPARs.....	35

Deploying an 8,000-user Microsoft Exchange Server 2010 Environment with Hitachi Compute Blade 2000 on the Hitachi Adaptable Modular Storage 2300

Implementation Guide

It can be difficult to design, deploy, manage and support Microsoft Exchange Server 2010 environments with hardware components from multiple vendors. This solution reduces the complexity of Exchange environments by using servers and storage from Hitachi that work together seamlessly.

The Hitachi Compute Blade 2000 used in this solution is an enterprise-class platform that offers the following features:

- Balanced system architecture that eliminates bottlenecks in performance and throughput
- Embedded Hitachi logical partitioning (LPAR) virtualization
- Unprecedented configuration flexibility
- Eco-friendly power-saving features and capabilities
- Fast recovery from server failures due to N+1 cold standby design that allows you to replace failed servers within minutes instead of hours or days

With its unique combination of power, efficiency and flexibility, you can now extend the benefits of virtualization to new areas of the enterprise data center — including mission-critical application servers and database servers like Exchange 2010 — while reducing total cost of ownership and maximizing simplicity..

The Hitachi Adaptable Modular Storage 2000 family is ideal for a demanding application like Exchange and delivers enterprise-class performance, capacity and functionality at a midrange price. It's a midrange storage product with symmetric active-active controllers that provide integrated, automated, hardware-based, front-to-back-end I/O load balancing.

This implementation guide focuses on deploying Exchange Server 2010 using the Hitachi Compute Blade 2000 with N+1 cold standby and logical partitioning technologies, the Hitachi Adaptable Modular Storage 2000 family and Exchange Server 2010's Database Availability Group (DAG) feature. It is based on the reference architecture described in the *Deploying an 8,000-user Exchange Server 2010 Environment with Hitachi Compute Blade 2000 and the Hitachi Adaptable Modular Storage 2300 Reference Architecture Guide* white paper. It is intended for use by IT administrators responsible for Exchange and storage. It assumes some familiarity with Hitachi Storage Navigator Module 2 software, Microsoft Windows Server 2008 R2 and Exchange Server 2010.

Tested Solution Components

The following sections detail the hardware and software components that were deployed in the Hitachi Data Systems lab.

Hardware Components

Table 2 lists the detailed information about the hardware components used in this solution.

Table 2. Hardware Components

<i>Hardware</i>	<i>Description</i>	<i>Version</i>	<i>Quantity</i>
Hitachi Adaptable Modular Storage 2300 storage system	Dual controller 16x 8GB Fibre Channel ports 16GB cache memory 80 SAS 2TB 7.2K RPM disks	0897/A-Y	1
Brocade 5300 switch	8Gb Fibre Channel ports	FOS 6.4.0E	2
Hitachi Compute Blade 2000 chassis	8-blade chassis 8 x 8GB dual-port HBAs 8 x 1GB network ports 2 x management modules 8 x cooling fan modules 4 x power supply modules	A0154-E-5234	1
Hitachi Compute Blade 2000 E55A2 blades	Full blade 2 x 6-Core Intel Xeon X5670 2.93GHz 80GB memory	58.22	4
Hitachi dual-port HBA	Dual port 8Gbps Fibre Channel PCIe card	4.2.6.670	8

Hitachi Compute Blade 2000 Chassis Configuration

This implementation guide uses four standard X55A2 blades, two standard 1Gb LAN switch modules, and eight Hitachi dual port 8Gbps HBA cards. Each blade has two on-board NIC, and each NIC is connected to a LAN switch module. Each blade has two PCIe slots available, and all are populated with Hitachi dual port 8Gbps HBA cards. Hitachi HBAs are required for HVM mode. With Hitachi HBAs, when HVM is enabled, eight virtual WWNs are created for LPARs. Figure 3 shows the front and back view of Hitachi Compute Blade 2000 used in this solution.




Figure 3

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

Table 3. Blade Servers

<i>Blade Server</i>	<i>LPAR</i>	<i>Server Name</i>	<i>Role</i>	<i>Number of CPU Cores</i>	<i>Memory (GB)</i>
Blade 0	N/A	N/A	N+1 failover standby	12	80
Blade 1	LPAR1	BS-Mgmt1	Blade management server	2	4
	LPAR2	BS-AD1	Active Directory DNS	4	8
Blade 2	LPAR1	BS-CASHT1	Client access Hub transport	6	16
	LPAR2	BS-MBX1	Mailbox server	6	64
Blade 3	LPAR1	BS-CASHT2	Client access Hub transport	6	16
	LPAR2	BS-MBX2	Mailbox server	6	64

This was the configuration of the blades for the environment built in the Hitachi Data Systems lab. The memory and processor cores allocated to the LPARs on blade 1 might vary in a production environment.

Software Components

Table 4 lists detailed information about the software components used in this solution.

Table 4. Software Components

<i>Software</i>	<i>Version</i>
Hitachi Storage Navigator Modular 2	Microcode dependent
Hitachi Dynamic Provisioning	Microcode dependent
Hitachi Dynamic Link Manager	6.5
Microsoft Windows Server	Windows 2008 R2 Enterprise
Microsoft Exchange Server 2010	SP1
Hitachi Server Conductor	80-90-A

Hitachi Dynamic Provisioning Software

On the Hitachi Adaptable Modular Storage 2000 family systems, Hitachi Dynamic Provisioning software provides wide striping and thin provisioning that dramatically improve performance, capacity utilization and management of your environment. By deploying your Exchange Server 2010 mailbox servers using volumes from Hitachi Dynamic Provisioning storage pools on the 2300, you can expect the following benefits:

- An improved I/O “buffer” to burst into during peak usage times
- A smoothing effect to the Exchange workload that can eliminate hot spots, resulting in reduced mailbox moves related to performance
- Minimization of excess, unused capacity by leveraging the combined capabilities of all disks comprising a storage pool

Hitachi Dynamic Link Manager Software

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

Hitachi Dynamic Link Manager software is installed on each server according to the instructions in the installation guide that accompanies the software.

Hitachi Server Conductor

Hitachi Server Conductor is suite of programs for centralized management of multiple servers. Server Conductor provides functions for managing servers efficiently, including functions to manage server software configurations and monitor server operating statuses and failures.

The Blade Server Manager Plus component of Server Conductor is required to implement an N+1 or N+M cold standby configuration.

Compared to a cluster configuration, an N+1 cold standby configuration reduces operating costs because a single standby server is shared by multiple active servers. Install Server Conductor on a server according to the instructions in the system administrator guide. The server must be rebooted when the installation completes.

This white paper is based on a lab environment that simulated a specific Exchange 2010 environment.

Solution Implementation

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

- USB DVD drive (comes with Compute Blades)
- Windows 2008 installation DVD or ISO image
- Driver kit version 12-01 CD or ISO image (includes driver for Hitachi HBA)
- 14 IP addresses:
 - One for the management module
 - Four for blades
 - Three for LPAR control (HVM)
 - Six for LPAR guest operating systems
- Four IP addresses in environments where these items are not already in place and connected:
 - Two for the storage system
 - One for each Fibre Channel switch

In addition, 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 the 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. SAN boot volumes for the OS are connected to two HBA ports. Exchange volumes are connected to four HBA ports. Four redundant paths from the switches to the 2300 are configured for Exchange volumes, and two redundant paths from the switches to the 2300 are used for SAN OS boot volumes.

Figure 4 shows the SAN connections that need to be made before implementing this solution.

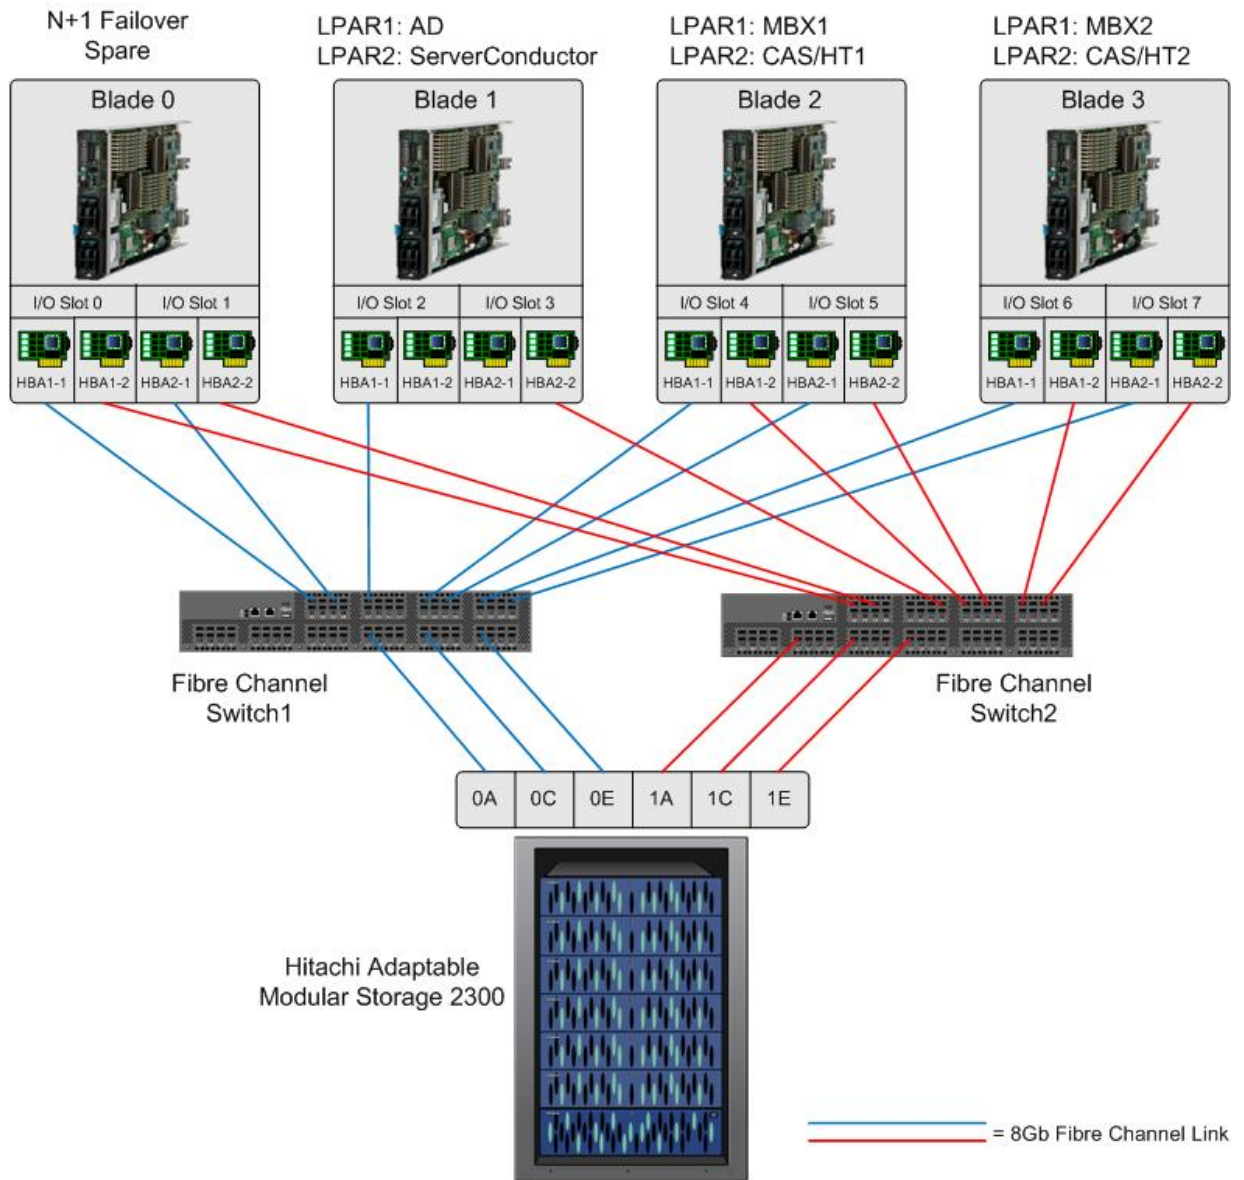


Figure 4

To deploy this Exchange Server 2010 solution, follow 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 LPARs.

6. Install Active Directory.
7. Install Hitachi Dynamic Link Manager software.
8. Install Exchange on LPARs.

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 HVM using a command-line interface. To access it, use terminal emulator client software to open 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 in window displays, click the **Launch Remote Console** button.

Hitachi is in the process of changing the product names for its blade products. Some of these interfaces might use old product names.

For more information about each of these high-level tasks, see the following documentation:

- Hitachi Storage Navigator Modular 2 online help
- Hitachi Dynamic Provisioning Software user's guide
- 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


The following sections describe how to configure the Hitachi Compute Blade 2000. These procedures assume that the blade chassis is racked, the blades are installed into the chassis, and that all cabling (LAN and SAN) is complete.

Configure Required Management IP Addresses

To configure the IP addresses for the management module and the blades using the management module web GUI, follow these steps:

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

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



2. Open a browser and enter `http://192.168.0.1/` in the URL field.

A log in screen displays.

3. Log in using the default user name `administrator` and the default password `password`.

The management module web GUI launches.

4. Click the **Settings** tab.

5. In the navigation tree on the left click **The configuration of network link**.

The **Management LAN network** pane displays.

6. Scroll down and click the **Edit** button in the **Management LAN network** pane.

The fields become editable.

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

You need one IP address for the management module and one IP address for each partition (blade).

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

The IP settings are saved.

9. Disconnect the Ethernet cable from the management module.
10. Connect an Ethernet cable from the management module to an external management LAN switch so you can perform management through any system console in the management LAN.

Configure Blades

The following procedures must be completed for each active blade.

Set Up HVM

This procedure assumes that you are connected to the management LAN. To set up HVM, follow these steps:

1. Open a browser and enter the blade IP address.
2. Log in to the blade web GUI using the following default user name **user01** and default password **pass01**.

The blade web GUI launches.

3. Click the **EFI Setup** link.

The **EFI Setup** pane displays.

4. In the **SMT (Simultaneous Multi-Threading)** field, select the **Disable** radio button.
5. In the **PCI Error Handling Mode** field, select the **Legacy** radio button.
6. Keep the default values for the other EFI settings and click the **Modify** button.

The **EFI Setup (Confirm)** pane displays.

7. Click **Confirm**.

The EFI settings are updated.

8. Click the **HVM Setup** link.

The **HVM Setup** pane displays.

9. In the **OS Mode** field, select the **HVM Mode** radio button.

10. Click **Modify**.

A confirmation window displays.

11. Click **Confirm**.

The window closes and the **HVM Setup** pane displays.

12. Click the **Power and LEDs** link.

The **Power and LEDs** pane displays.

13. Click the **Power on** button.

The blade is powered on.

Configure HVM

To configure HVM, follow these steps:

1. Initiate a Telnet session to the IP address assigned to the partition.
2. Enter your user name and password.

If you have not changed these from the defaults, use the default user name **adminstrator** and the default password **password**.

The **System Configuration** window displays with a message that indicates that HVM is being initialized.

3. When the message clears, place your cursor on **HVM IP Address** and press the Enter key.

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

4. Enter an IP address and press the Enter key.
5. Place your cursor on **Subnet Mask** and press the Enter key.

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

6. Enter a subnet mask and press the Enter key.
7. Place your cursor on **Default Gateway** and press the Enter key.

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

8. Enter the IP address for the default gateway and press the Enter key.
9. Place your cursor on **VNIC System No** and press the Enter key.

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

10. Type a number and press the Enter key.

The VNIC system number is used to create a MAC address that is unique and does not conflict with other MAC addresses. Specify a unique value between 1 and 128 for each HVM.

11. Press the F10 key.

A confirmation menu displays.

12. Place your cursor on **Yes** and press the Enter key.

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

13. Press the Enter key.

The window closes.

14. Press the Escape key.

15. The main menu displays.

Configure an LPAR

Configuring an LPAR involves assigning resources, such as CPU, memory, network and Fibre Channel, and so on, to it. This procedure uses the Telnet session launched in the previous procedure. This procedure must be completed for each LPAR. To set up an LPAR, follow these steps:

1. Place your cursor on **Logical Partition Configuration** and press the Enter key.

The **Logical Partition (LPAR) Configuration** menu displays.

2. Press the F6 key.

The **Add LPAR** menu displays.

3. Place your cursor on the first **NO_NAME** and press the Enter key.

The first line of the **Logical Partition (LPAR) Configuration** pane is populated.

4. Place your cursor on **NO_NAME** and press the Enter key.

The **Logical Partition Name** window displays.

5. Enter a name for the logical partition and press the Enter key.

6. Place your cursor on the **Shr** column for shared CPU mode, and press the Enter key

A **The number of Shared Logical Processors** window displays.

7. Type the number of CPUs for the LPAR and press the Enter key.

- Place your cursor on the **Mem** column and press the Enter key.

A **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        |                                     |         ||
|| 6        |                                     |         ||
|| 7        |                                     |         ||
|| 8        |                                     |         ||
|| 9        |                                     |         ||
|| 10       |                                     |         ||
||                                     | Maximum size : 4608 |p / [PageDown]:Page Down ||
+-----+| Effectitive max : 4608 |-----+
|+ Logical Information -----+ Physical Information -----+
||                                     |er Memory : 4608 MB || |
|| Assign Total   5  0+-----+ocessors : 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 Dsp F3:Act F4:Deact F5:React F6:Add F7:Remove Esc:Menu |
+-----+

```

- Press the F1 key.

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

- Type the memory size in GB and press the Enter key.

The menu closes.

- Place your cursor on an asterisk (*) in **AA** column and press the Enter key.

An auto activation order menu displays.

When you set auto activate on LPARs, the LPARs are automatically powered on in the order that you specify after you power on the HVM.

- Type the order number and press the Enter key.

The menu closes.

- Press the Escape key.

The main menu displays.

- Place your cursor on **PCI Device Assignment** and press the Enter key.

The **PCI Device Assignment** window displays.

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

The **PCI Device Scheduling mode Assignment** menu displays. This is where you specify that you want to use 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 LPAR1  Dea   A  -  -  -  -  -
|| 2 LPAR2  Dea   |-----|
|| 3 LPAR3  Dea   | PCI Device Scheduling mode Assignment |
|| 4 LPAR4  Dea   |
|| 5        |
|| 6        |           D
|| 7        |-----|
|| 8        | PCI Device is Shared Mode
|| 9        |-----|
|| 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 |
+-----+

```

- Place your cursor on **S** for shared mode and press the Enter key.

The menu closes.

Perform Step 15 and Step 16 for every PCI device.

- Press the F10 key to save the settings.

The **Save Setting** window displays.

- Place your cursor on Yes and press the Enter key.

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

- Press the Enter key.

The message closes.

20. Press the Escape key.

The main menu displays.

21. Place your cursor on **VNIC Assignment** and press the Enter key.

The **Virtual NIC Assignment** window displays.

22. Put cursor on the asterisk (*) below **Virtual NIC Number 0** and press the Enter key.

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

23. Place your cursor on **1a** and press the Enter key.

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

24. Press the Escape key.

The main menu displays.

25. Place your cursor on **Shared FC Assignment** and press the Enter key.

The **Shared FC Assignment** menu displays.

26. Place your cursor on an asterisk (*) and press the Enter key.

The **Shared FC Assignment** menu displays.

```
Shared FC Assignment
|| Shared FC#: 0 1 2 3 4 5 6 7 8 9 10 11 || | | | | | | | | | | |
|| Slot#: 0 0 ||
|| Port#: 0 1 ||
|| PortStatus: ||
|| # Name Sta | Shared FC vfcWwnId Assignment | ||
|| 1 LPAR1 Dea | | | | | | | | | | | ||
|| 2 LPAR2 Dea | | | | | | | | | | | ||
|| 3 LPAR3 Dea | | | | | | | | | | | ||
|| 4 LPAR4 Dea | | | | | | | | | | | ||
|| 5 | | | | | | | | | | | ||
|| 6 | | | | | | | | | | | ||
|| 7 | | | | | | | | | | | ||
|| 8 | | | | | | | | | | | ||
|| 9 | | | | | | | | | | | ||
|| 10 | | | | | | | | | | | ||
|| [PageUp]:Page Up / [PageDown]:Page Down ||
|| Selected Virtual FC Port Wwn Information: ||
|| # LPAR# Wwn Wwn Bus# Dev# Func# vfcID# ||
|| 0 1 0 0 30 4 0 -- ||
|| F11:Left F12:Right Esc:Menu ||
```

27. Place your cursor on the first number available on the list and press the Enter key.

Repeat Step 26 and Step 27 for all shared Fibre Channel numbers.

Configure the SAN

Port 0A, 1A, 0E and 1E on the Hitachi Adaptable Modular Storage 2300 are used for Exchange databases and logs, and port 0C and 1C are used for SAN OS boot volumes. Hitachi Dynamic Link Manager software is used for multipathing with the round-robin load balancing algorithm.

The following sections describe how to configure the SAN for this solution. Some of the configuration details are specific to the Hitachi Data Systems lab environment and might vary in a production environment.

Zone SAN Switches

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

Configure two zones for each host (SAN OS boot). Each zone must use different switches to provide redundancy. Configure four zones for each 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. Enable the zones.

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

Table 5 lists the zoning deployed in the Hitachi Data Systems lab.

Table 5. Zoning Configuration

<i>Blade Number</i>	<i>LPAR Host</i>	<i>Host HBA Number</i>	<i>Storage System Host Group</i>	<i>Storage System Port</i>	<i>Zone Name</i>
Blade 0	N/A	HBA1_1	N/A	N/A	N/A
		HBA1_2	N/A	N/A	N/A
		HBA2_1	N/A	N/A	N/A
		HBA2_2	N/A	N/A	N/A
Blade 1	LPAR1	HBA1_1	B1L1_HBA1_1	0 C	B1L1_HBA1_1_0C
		HBA1_2	B1L1_HBA1_2	1 C	B1L1_HBA1_2_1C
		HBA1_1	B1L1_HBA1_1	0 A	B1L1_HBA1_1_0A
		HBA1_2	B1L1_HBA1_2	1 A	B1L1_HBA1_2_1A

		HBA2_1	B1L1_HBA2_1	0E	B1L1_HBA2_1_0E	
		HBA2_2	B1L1_HBA2_2	1E	B1L1_HBA2_2_1E	
	LPAR2	HBA1_1	B1L2_HBA1_1	0 C	B1L2_HBA1_1_0C	
		HBA1_2	B1L2_HBA1_2	1 C	B1L2_HBA1_2_1C	
		HBA1_1	B1L2_HBA1_1	0 A	B1L2_HBA1_1_0A	
		HBA1_2	B1L2_HBA1_2	1 A	B1L2_HBA1_2_1A	
		HBA2_1	B1L2_HBA2_1	0E	B1L2_HBA2_1_0E	
		HBA2_2	B1L2_HBA2_2	1E	B1L2_HBA2_2_1E	
Blade 2	LPAR1	HBA1_1	B2L1_HBA1_1	0 C	B2L1_HBA1_1_0C	
		HBA1_2	B2L1_HBA1_2	1 C	B2L1_HBA1_2_1C	
		HBA1_1	B2L1_HBA1_1	0 A	B2L1_HBA1_1_0A	
		HBA1_2	B2L1_HBA1_2	1 A	B2L1_HBA1_2_1A	
		HBA2_1	B2L1_HBA2_1	0E	B2L1_HBA2_1_0E	
		HBA2_2	B2L1_HBA2_2	1E	B2L1_HBA2_2_1E	
	LPAR2	HBA1_1	B2L2_HBA1_1	0 C	B2L2_HBA1_1_0C	
		HBA1_2	B2L2_HBA1_2	1 C	B2L2_HBA1_2_1C	
		HBA1_1	B2L2_HBA1_1	0 A	B2L2_HBA1_1_0A	
		HBA1_2	B2L2_HBA1_2	1 A	B2L2_HBA1_2_1A	
		HBA2_1	B2L2_HBA2_1	0E	B2L2_HBA2_1_0E	
		HBA2_2	B2L2_HBA2_2	1E	B2L2_HBA2_2_1E	
	Blade 3	LPAR1	HBA1_1	B3L1_HBA1_1	0 C	B3L1_HBA1_1_0C
			HBA1_2	B3L1_HBA1_2	1 C	B3L1_HBA1_2_1C
HBA1_1			B3L1_HBA1_1	0 A	B3L1_HBA1_1_0A	
HBA1_2			B3L1_HBA1_2	1 A	B3L1_HBA1_2_1A	
HBA2_1			B3L1_HBA2_1	0E	B3L1_HBA2_1_0E	
HBA2_2			B3L1_HBA2_2	1E	B3L1_HBA2_2_1E	
LPAR2		HBA1_1	B3L2_HBA1_1	0 C	B3L2_HBA1_1_0C	
		HBA1_2	B3L2_HBA1_2	1 C	B3L2_HBA1_2_1C	
		HBA1_1	B3L2_HBA1_1	0 A	B3L2_HBA1_1_0A	
		HBA1_2	B3L2_HBA1_2	1 A	B3L2_HBA1_2_1A	
		HBA2_1	B3L2_HBA2_1	0E	B3L2_HBA2_1_0E	
		HBA2_2	B3L2_HBA2_2	1E	B3L2_HBA2_2_1E	

Configure Storage

The following sections describe storage configuration steps for this solution. These procedures assume that all appropriate licenses are installed on your storage system.

Configure Fibre Channel Port Settings

In preparation for this step, all the HBAs on Exchange must be connected to a Fibre Channel switch and appropriately zoned to the storage system ports.

To ensure that your storage Fibre Channel ports are configured for the correct topology configuration using Hitachi Storage Navigator Modular 2 software, follow these steps:

1. Log in to Storage Navigator Modular 2 software.

Ensure that you have modify privileges on Storage Navigator Modular 2 software.

2. Click the **Array Name** link to open the storage system to be used for the Exchange Server environment.

You might be prompted for your user name and password for the storage system.

3. Expand the **Settings** heading and click the **FC Settings** link.
4. Notice the **Topology** column for the ports that were zoned to connect to the Exchange hosts on the SAN.

If any of the ports are not configured as Point-to-Point, you must change their configuration.

5. Click the **Edit FC Port** button.
6. Choose **Point-to-Point** from the **Topology** drop-down menu and click **OK**.

A disclaimer displays indicating that the change will interrupt I/O between any host that is connected to the port at the moment.

7. Click the **Confirm** button and wait a few seconds for the change to take place.

After the connection between the storage system and the host is established, the **FC Settings** window shows all ports with a **LinkUp(F_Port Connected)** status.

Enable Host Group Security

Enabling host group security allows you to create multiple host groups on a single storage port, which in turn allows you to isolate traffic for an HBA.

To enable host group security using Hitachi Storage Navigator Modular 2 software, follow these steps:

1. Log in to Storage Navigator Modular 2 software.

Ensure that you have modify privileges on Storage Navigator Modular 2 software.

2. Click the **Array Name** link to open the storage system to be used for the Exchange Server environment.
3. Expand the **Groups** heading and click the **Host Groups** link.
4. Click the **Host Group Security** tab and ensure that port security is enabled on all ports.

If the host security is not enabled, select all ports to be used in the environment and click the **Change Host Group Security** button.

5. Select the **Yes** checkbox and click **OK**.

The **Host Security Group** setting is now enabled on all selected ports.

Configure Storage for SAN OS Boot

This white paper provides step-by-step instructions explaining how to provision and manage the storage for your SAN OS boot environment using the Storage Navigator Modular 2 software GUI.

This solution uses RAID-5 (4D+1P) for SAN OS boot.

Table 6 lists details about the 100GB LUs created for SAN OS boot and mapped to storage ports 0C and 1C.

Table 6. LPAR Host LUs

<i>LPAR Host</i>	<i>LU Number</i>
BS-Mgmt1	0012
BS-AD1	0013
BS-MBX1	0014
BS-CASHT1	0015
BS-MBX2	0016
BS-CASHT2	0017

Create RAID Groups

To create RAID Groups using Hitachi Storage Navigator Modular 2 software, follow these steps:

1. Log in to Storage Navigator Modular 2 software.

Ensure that you have modify privileges on Storage Navigator Modular 2 software.

2. Click the **Array Name** link to open the storage system to be used for the Exchange Server environment.
3. Expand the **Groups** heading in the storage system's pane and click the **Logical Units** link.

The right pane has three tabs: **Logical Units**, **RAID Groups** and **DP Pools**.

4. Click the **RAID Groups** tab and click the **Create RG** button.
5. Select **RAID5** from the **RAID Level** drop-down menu and **4D+1P** from the **Combination** drop-down menu.
6. Select the **Automatic Selection** radio button.

Enabling automatic selection means that Storage Navigator Modular 2 software selects the next available drives of the type shown in the box. Hitachi Data Systems recommends enabling this setting.

7. Click **OK**.

A message displays stating that the RAID group was successfully created.

Configure Storage for Exchange

This solution uses RAID-1+0 (2D+2D) for both database and transaction logs. Exchange databases are stored in DP pool 0 and DP pool 1 and logs are stored in DP pool 2 and DP pool 3. Table 7 lists details for the Exchange database and log used in the Dynamic Provisioning pools.

Table 7. Exchange Database and Log Dynamic Provisioning Pools

<i>Dynamic Provisioning Pool</i>	<i>Number of RAID Groups</i>	<i>Number of Drives</i>	<i>Usable Pool Capacity (TB)</i>
0	9	36	30.0
1	9	36	30.0
2	1	4	3.5
3	1	4	3.5

This solution uses 24 LUs for the active and passive copies of the Exchange databases and logs and mapped to storage ports 0A, 1A, 0E and 1E. Table 8 lists details of LUs used in this lab.

Table 8. Exchange Database and Log LUs

<i>Host</i>	<i>Purpose</i>	<i>LU</i>	<i>Size (GB)</i>	<i>Dynamic Provisioning Pool Number</i>
BS-MBX1	Database	0000	2000	000
BS-MBX1	Database	0001	2000	000
BS-MBX1	Database	0002	2000	000
BS-MBX1	Database copy	0003	2000	000
BS-MBX1	Database copy	0004	2000	000
BS-MBX1	Database copy	0005	200	000
BS-MBX1	Log	0020	100	002
BS-MBX1	Log	0021	100	002
BS-MBX1	Log	0022	100	002
BS-MBX1	Log copy	0023	100	002
BS-MBX1	Log copy	0024	100	002
BS-MBX1	Log copy	0025	100	002
BS-MBX2	Database	0006	2000	001
BS-MBX2	Database	0007	2000	001
BS-MBX2	Database	0008	2000	001
BS-MBX2	Database copy	0009	2000	001
BS-MBX2	Database copy	0010	2000	001
BS-MBX2	Database copy	0011	2000	001
BS-MBX2	Log	0026	100	003
BS-MBX2	Log	0027	100	003
BS-MBX2	Log	0028	100	003
BS-MBX2	Log copy	0029	100	003
BS-MBX2	Log copy	0030	100	003
BS-MBX2	Log copy	0031	100	003

Create DP Pools

To create a DP pool using Hitachi Storage Navigator Modular 2 software, follow these steps:

1. Log in to the Storage Navigator Modular 2 instance in your data center.

You must have modify privileges for Storage Navigator Modular 2 software.

2. Click the **Array Name** link to open the storage system to be used for the Exchange Server environment.
3. Expand the **Groups** heading in the storage system's pane window and click the **Logical Units** link.

The right pane displayed has three tabs: **Logical Units**, **RAID Groups** and **DP Pools**.

4. Click the **DP Pools** tab and click the **Create Pool** button.

The **Create DP Pool** window displays.

5. Select **RAID1+0** from the **RAID Level** drop-down menu and **2D+2D** from the **Combination** drop-down menu.

Enter the information for the DP pool to be created.

Basic | Advanced

* DP Pool:
From 0 to max (array model type dependent)

RAID Level:

Combination:

* Number of drives:
From 2 to max (based on drive count)

Drives:

Automatic Selection:

Drive Type:

Drive Capacity:

Manual Selection:

Tray	HDU	Drive Type	Status
01	14	SAS7K (2TB)	Mounted

* Required field

OK Cancel

Done

The **Number of drives** field automatically changes based on your RAID level and combination choices.

Note that the pool consists of a single RAID group at creation time, and after a RAID level and combination are selected, all future RAID groups added to the Dynamic Provisioning pool also have the same settings.

6. Select the **Automatic Selection** radio button.

If you have different types of drives installed in the storage system (for example SATA and SAS, or different capacities), be sure to select the correct type so that Storage Navigator Modular 2 software can select the right type of hard drive.

Enabling automatic selection means that Storage Navigator Modular 2 software selects the next available drives of the type shown in the box. Hitachi Data Systems recommends enabling this setting.

7. Click the **Advanced** tab, modify any of the settings according to your environment requirements and click **OK**.

A message displays indicating that the pool was successfully created.

8. Click the **Close** button.

Note that the pool immediately starts the formatting process in the background.

9. To increase the pool capacity, click the **DP Pools** tab, select the pool's check box and click the **Add Pool Capacity** button.

10. Select the **Automatic Selection** radio button and click **OK**.

Enabling automatic selection means that Storage Navigator Modular 2 software selects the next available drives of the type shown in the box. Hitachi Data Systems recommends using this setting.

11. Repeat these steps for DP pool 1, DP pool 2 and DP pool 3.

Create LUs

This procedure creates LUs for the operating system and Exchange. To create LUs using Hitachi Storage Navigator Modular 2 software, follow these steps:

1. Log in to Storage Navigator Modular 2 software.

Ensure that you have modify privileges on Storage Navigator Modular 2 software.

2. Click the **Array Name** link to open the storage system to be used for the Exchange Server environment.
3. Expand the **Groups** heading in the storage systems' pane window and click the **Logical Units** link.

4. Click the **Create LU** button.

The **Create Logical Unit** pop-up window displays.

5. In the **Type** field, Select the **RAID Group** radio button for the SAN boot LU.
6. From the **RAID Group/DP Pool Number** drop-down menu, select **000**.
7. In the **LUN** field, type a logical unit number.

This information is listed in Table 6.

8. In the **Capacity** field, enter **100** and select **GB** from the drop-down menu.
9. Leave the Accelerate Wide Striping Mode checkbox unchecked.
10. Click the **OK** button.

The **Create Logical Unit** pane refreshes, populated with the new LUN information.

11. Repeat Step 4 through Step 10 inclusive for every SAN boot LU.
12. Click the **Create LU** button.

The **Create Logical Unit** pop-up window displays.

13. In the **Type** field, Select the **DP Pool** radio button for the SAN boot LU.
14. From the **RAID Group/DP Pool Number** drop-down menu, select a DP pool number.

This information is listed in Table 8.

15. In the **LUN** field, type a logical unit number.

This information is listed in Table 8.

16. In the **Capacity** field, type a capacity amount and select **GB** from the drop-down menu.

This information is listed in Table 8.

17. Leave the **Accelerate Wide Striping Mode** checkbox unchecked.
18. Click the **OK** button.

The **Create Logical Unit** pane refreshes, populated with the new LUN information.

19. Repeat Step 12 through Step 18 inclusive for every Exchange log and database LU.

Create Host Groups

To create host groups and assign the SAN Boot OS LUs to the host groups using Hitachi Storage Navigator Modular 2 software, follow these steps:

1. Log in to Storage Navigator Modular 2 software.

Ensure that you have modify privileges on Storage Navigator Modular 2 software.

2. Click the **Array Name** link to open the storage system to be used for the Exchange Server environment.
3. Expand the **Groups** heading and click the **Host Groups** link.
4. Click the **Create Host Group** button.

The **Host Group No.** field can remain the same because Storage Navigator Modular 2 software automatically selects the next available number.

5. Enter a host group name in the **Name** field.

Use a name that identifies the hosts connected to it.

6. Select port 0C and port 1C for the SAN OS boot hosts.
7. This ensures that multiple paths to the LUs to be used are available both for high availability and failover purposes.

Note that if the HBA ports are not detected by the storage system, their WWNs are not shown in the **Detected WWNs** window. In that case, it might be necessary to troubleshoot the environment to ensure that the proper settings are set on the storage ports, host HBAs and switches.

8. Highlight the World Wide Names to be added to the host group and click **Add**.

The window automatically refreshes and the added WWNs display in the **Selected WWNs** column.

9. Click the **Logical Units** tab.
10. Select the LUs to be added to the host group and click **Add**.

The screen automatically refreshes and the added LUs display in the **Assigned Logical Units** column.

11. Click the **Options** tab.
12. Select **Windows** from the **Platform** drop-down menu and click **OK**.
13. Repeat this procedure for the rest of the SAN OS boot hosts and Exchange databases and logs.

Set Up SAN OS Boot

SAN boot is the capability to boot the OS from a LUN on an Adaptable Modular Storage 2300 without using any of the server's internal disks. SAN boot is required for LPARs and the N+1 cold standby feature. The following sections describe how to set up the SAN OS boot.

These procedures must be repeated for each LPAR.

Assign USB Devices and the Remote Console to an LPAR

Skip this procedure when setting up the first LPAR. By default, the USB devices and remote console are assigned to the first LPAR.

USB ports and the remote console can be shared between active LPARs. However, they can be assigned to only one LPAR at a time. Before you set up a second LPAR on a blade, you must first assign USB devices and the remote console to the second LPAR.

To assign USB devices and the remote console to an LPAR, follow these steps:

1. Launch the blade server command-line interface.
2. On the main menu, place your cursor on **PCI Device Assignment** and press the Enter key.

The **PCI Device Assignment** window displays.

3. Press the F5 key.

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

4. Place your cursor on the target LPAR that you want to attach and press the Enter key.

A confirmation message displays.

5. Place your cursor on **Yes** and press the Enter key.

A message indicating that the change is being made displays.

The PCI device is attached to the selected LPAR.

Activate an LPAR

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

To activate an LPAR using the blade management command line interface, follow these steps:

1. On the main menu, press the F3 key.

The **Activate LPAR** menu displays.

2. Place your cursor on **an LPAR to be activated** and press the Enter key.
3. Place your cursor on **Continue (Don't show this message)** and press the Enter key.

The selected LPAR is powered on.

Configure HBAs

This procedure assumes that the Telnet session from the previous procedure is still open.

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

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

A log in screen displays.

2. Click the **Launch Remote Console** button.

The remote console log in window appears.

3. Enter the default User ID `user01` and default password `pass01`.

The remote console launches.

4. Place your cursor on **Continue** and press the Enter key.

The EFI shell launches.

5. Execute the `drivers` command.

A list of drivers displays.

6. Look for **Hitachi PCI-X/PCIe Fibre channel DR** in the list of driver names and note the driver number in the left column labelled **DRV**.

In this case, 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
.
```

7. Execute the `drvcfg <driver number>` command.

In this case, the command is `drvcfg 70`.

A list of configurable components displays.

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

8. Note the driver number and controller number of the first configurable component on the list.

9. Execute the `drvcfg -s <driver number> <controller number>` command.

In this case, the command is `drvcfg -s 70 7B`.

The command line prompt changes to `hfccfg>`.

10. Execute a `select` command.

A list of Fibre Channel ports and WWNs displays.

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

The command-line prompt changes to `hfccfg.<portwwn>>`.

This solution uses Fibre Channel port #3.

12. Execute the `set` command.

A configuration dialog launches.

13. Ensure 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

14. Configure the first WWN in the boot device list by following these steps:

- a) At the **Change (y/n)** prompt, enter y.

A **Please select a number** prompt displays.

- b) Enter 1 to select a WWN to configure.

A list of options displays.

- c) Enter 1 to scan for targets.

A list of target devices displays.

- d) Enter 1 to scan for LUNs.

A list of LUNs displays.

e) Enter 1 to choose a LUN.

A list showing the target device's WWN and LUN displays.

An **Update list#1 (y/n)** prompt displays.

f) Enter y.

The updated boot device list displays and a message indicating that the base settings are complete displays.

You are returned to the **hfccfg.<wwn>>** prompt.

15. Execute a save command.

A confirmation prompt displays.

16. Enter y to confirm the save command.

A confirmation message displays.

The path to the storage for the boot OS is configured.

The **hfccfg<wwn>>** prompt displays.

17. Execute the `exit` command.

You are returned to the main EFI shell.

Map OS Boot Device

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

1. Execute the `reconnect -r` command.

A success message displays.

2. Execute the `map -r` command to remap the devices.

A device mapping table displays.

3. Execute the `exit` command.

You are returned to the EFI main menu.

4. Place your cursor on **Boot Maintenance Manager** and press the Enter key.


The **Boot Maintenance Manager** window displays.

5. Place your cursor on **Boot Options** and press the Enter key.

The Boot Maintenance Manager Boot Options window displays.

6. Place your cursor on **Add Boot Option** and press the Enter key.

The File Explorer window displays.

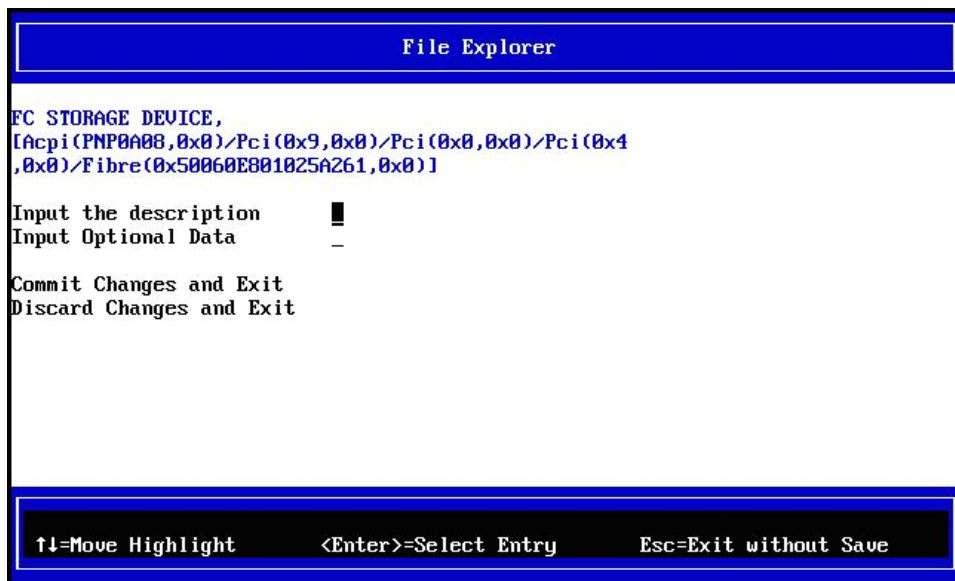


7. Place your cursor on **FC STORAGE DEVICE** and press the Enter key.

Multiple input fields display.

8. Place your cursor on the **Input the description** field and press the Enter key.

The field becomes editable.



9. Type a name for the device and press the Enter key.

The prompt closes.

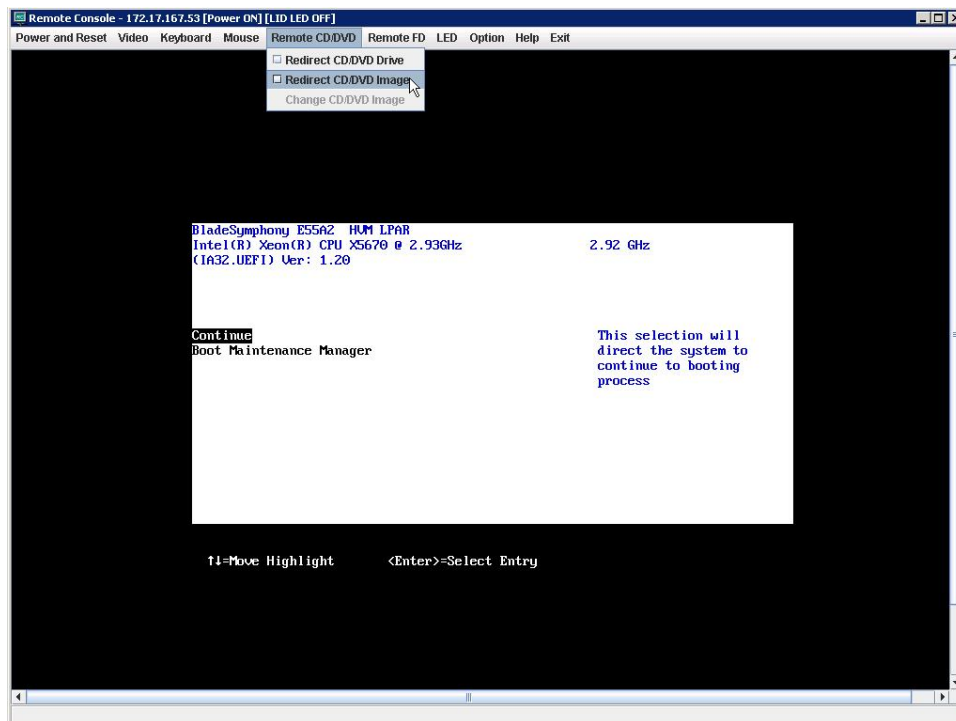
10. Place your cursor on **Commit Changes and Exit** and press the Enter key.

You are returned to the EFI main menu.

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

This procedure assumes that the remote console session from the “Configure HBAs” procedure is still open. To mount the Windows 2008 R2 ISO image from a remote DVD, follow these steps:

1. From the menu bar, select **Remote CD/DVD > Redirect CD/DVD Image**.



The **Open** window displays.

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

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

3. Place your cursor on **Boot Maintenance Manager** and press the Enter key.

The **Boot Maintenance Manager** window displays.

4. Place your cursor on **Boot Options** and press the Enter key.

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

5. Place your cursor on **Add Boot Option** and press the Enter key.
6. The **File Explorer** window displays.
7. Ensure that the USB storage device is listed in the **File Explorer** window.

If the USB storage device is not listed, reactivate the LPAR and try again. To reactivate the LPAR, go to the blade server command-line interface and on the main menu, press the F5 key to select a LPAR to be reactivated. Soon after the reactivation, the EFI main screen displays on remote console. Return to Step 3 and continue.

8. Place your cursor on **USB STORAGE DEVICE** and press the Enter key.

Multiple input fields display.

9. Place your cursor on the **Input the description** field and press the Enter key.

The field becomes editable.

10. Type a name for the device and press the Enter key.

The prompt closes.

11. Place your cursor on **Commit Changes and Exit** and press the Enter key.

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

12. Place your cursor on **Change Boot Order** and press the Enter key.

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

13. Place your cursor on **Change the order** and press the Enter key.

The **Boot Order** menu displays.

14. Use the plus (+) and minus (–) keys to change the boot order to use the remote DVD first and the Fibre Channel storage LUN second, and press the Enter key.

15. Place your cursor on **Commit Change and Exit** and press the Enter key.

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

16. Press the Escape key.

You are returned to the main menu.

17. Place your cursor on **Continue** and press the Enter key.

The LPAR begins to boot from the Windows ISO image.

Install Windows 2008 R2 on an LPAR

After the LPAR boots from the ISO image, follow Microsoft's instructions to complete the installation, specifying the Hitachi X55X2 HBA driver, Windows Server 2008 R2 Enterprise Edition (full installation) and a strong password for the local Administrator account. You must install Windows 2008 R2 on all active LPARs. Ensure that all Windows patches are up to date.

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

Install Active Directory

For this solution, a domain controller was installed on LPAR 1 of blade 1. In a production environment, add this domain controller to the existing Active Directory infrastructure. Configure this domain controller as a global catalog.

To install Active Directory on a server, follow Microsoft's instructions. See the Microsoft TechNet article "[Install Active Directory on a Member Server.](#)"

After you configure Active Directory, you must add the Windows servers to the domain.

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 LPARs

Install the Exchange mailbox role on LPAR 2 on blade 2 and LPAR 2 on blade 3. Install the client access and hub transport Exchange roles on LPAR 1 on blade 2 and LPAR 1 on blade 3.

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

Table 9 shows the database and log drive letter layout used in the Hitachi Data Systems lab.

Table 9. Exchange Database and Log LUNs

<i>Mailbox Server</i>	<i>Database Name</i>	<i>Database Drive Letter</i>	<i>Log Name</i>	<i>Log Drive Letter</i>
BS-MBX1	DB1 (active)	H	LOG1 (active)	N
	DB2 (active)	I	LOG2 (active)	O
	DB3 (active)	J	LOG3 (active)	P
	DB4 (passive)	E	LOG4 (passive)	K
	DB5 (passive)	F	LOG5 (passive)	L
	DB6 (passive)	G	LOG6 (passive)	M
BS-MBX2	DB1 (passive)	H	LOG1 (passive)	N
	DB2 (passive)	I	LOG2 (passive)	O
	DB3 (passive)	J	LOG3 (passive)	P
	DB4 (active)	E	LOG4 (active)	K
	DB5 (active)	F	LOG5 (active)	L
	DB6 (active)	G	LOG6 (active)	M

Create the DAG

You must create the DAG after the Exchange mailbox servers are installed and configured. 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.](#)"

 **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-075-01 February 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
