

Deploying Microsoft® SQL Server® 2008 R2 with Logical Partitioning on the Hitachi Virtual Storage Platform with Hitachi Dynamic Tiering

Implementation Guide

By Eduardo Freitas

July 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

Solution Overview	3
Key Solution Components.....	5
Hitachi Virtual Storage Platform	5
Hitachi Dynamic Tiering	5
Hitachi Compute Blade 2000.....	6
Microsoft SQL Server 2008 R2	6
Hitachi Dynamic Link Manager Advanced	7
Tested Solution Components.....	7
Hardware Components	7
Storage Area Network Components	10
Software Components.....	11
Solution Implementation.....	11
Implementation Preparations	12
Configure the Hitachi Compute Blade 2000.....	13
Configure the SAN	20
Configure the Virtual Storage Platform	23
Set Up SAN OS Boot	27
Install Windows 2008 R2 on an LPAR	36
Install Hitachi Dynamic Link Manager	36
Install Microsoft SQL Server 2008 R2.....	36
Configure Microsoft SQL Server 2008 R2.....	37

Deploying Microsoft® SQL Server® 2008 R2 with Logical Partitioning on the Hitachi Virtual Storage Platform with Hitachi Dynamic Tiering

Implementation Guide

IT administrators need database solutions that provide data protection and simple management in environmentally friendly data centers. Using Microsoft SQL Server 2008 R2 with SAN storage, like the Hitachi Virtual Storage Platform, accomplishes those objectives. This solution reduces the complexity of SQL server environments by using servers and storage from Hitachi that work together seamlessly to lower administration costs, improve system efficiency and capacity utilization, and enable virtualization.

Microsoft SQL Server 2008 R2 can be deployed for many types and numbers of users in a wide variety of infrastructure topologies. By implementing it on the Hitachi Compute Blade using the Hitachi Virtual Storage Platform with Hitachi Dynamic Tiering, you can scale from a small environment to a very large environment.

This implementation guide focuses on deploying Microsoft SQL Server 2008. It is based on the reference architecture described in [Deploying Microsoft SQL Server 2008 R2 with Microsoft Hyper-V on the Hitachi Virtual Storage Platform with Hitachi Dynamic Tiering Reference Architecture Guide](#). It is intended for use by IT administrators responsible for managing Microsoft SQL Server installations, servers, and storage. It assumes some familiarity with Hitachi Storage Navigator software, Microsoft Windows Server 2008 R2 and SQL Server 2008.

Solution Overview

This describes how to deploy a Microsoft SQL Server 2008 R2 environment that supports 75,000 users on a Hitachi Virtual Storage Platform storage system using Hitachi Dynamic Tiering. For ease of management and scalability, and to provide predictable performance, this solution uses a building block approach. This building block was duplicated, along with adding a second virtual machine and SQL server instance, to support 150,000 users.

A building block consists of a single logical partition running the following:

- Microsoft Windows Server 2008 R2
- Microsoft SQL Server 2008 R2
- Underlying storage from the Hitachi Virtual Storage Platform.

The building block for the 75,000-user account database consists of the following:

- SAS and SSD drives, using Hitachi Dynamic Tiering
- LUNs provisioned from dynamic tiering pools on the Hitachi Virtual Storage Platform:
 - Two 450GB LUNs for the SQL database
 - One 90GB LUN for tempdb
 - One 225GB LUN for the SQL transaction log

Figure 1 shows the high level design of this solution.




Figure 1

Key Solution Components

The following describes the key components needed to deploy this solution.

Hitachi Virtual Storage Platform

The Hitachi Virtual Storage Platform is a 3D scaling storage platform. With the unique ability to scale up, scale out, and scale deep at the same time in a single storage system, the Virtual Storage Platform flexibly adapts for performance, capacity, connectivity, and virtualization.

- **Scale Up**—Increase performance, capacity, and connectivity by adding cache, processors, connections, and disks to the base system.
- **Scale Out**—Combine multiple chassis into a single logical system with shared resources.
- **Scale Deep**—Extend the advanced functions of the Virtual Storage Platform to external multivendor storage.

For more information, see the [Hitachi Virtual Storage Platform](#) on the Hitachi Data Systems website.

Hitachi Dynamic Tiering

Hitachi Dynamic Tiering eliminates manual data classification and movement between storage tiers. This optimizes tiered storage usage while improving performance.

Instead of manually provisioning space from several storage technologies with different performance and cost characteristics, Hitachi Dynamic Tiering enables the management of multiple storage tiers as a single entity. By leveraging the thin provisioning and wide striping features of Hitachi Dynamic Provisioning, Hitachi Dynamic Tiering presents a virtual volume with embedded smart tiering. It monitors access and moves data at the 42MB page level.

Breaking the volume into pages, 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.

After an initial setup process, Hitachi Dynamic Tiering monitors data access in real time. It makes decisions on moving data between the available storage tiers based on actual use. Using this approach, Hitachi Dynamic Tiering improves the availability and performance of your storage systems and the applications using that storage.

Hitachi Dynamic Tiering on the Virtual Storage Platform allows a single pool to contain tiers made up of differently-arranged RAID groups using any type of disk. It manages data migration between the various tiers within a pool automatically. This eliminates most user management of storage tiers within a storage system, and maintains peak performance under dynamic conditions without storage administrator intervention.

Hitachi Dynamic Tiering provides virtual storage capacity to eliminate application service interruptions, to reduce costs, and to simplify administration, as follows:

- Optimize storage performance and capacity based on business or application requirements.
- Support deferring storage capacity upgrades until required by actual usage.
- Simplify the processes used for administration.
- 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.

For more information, see [Hitachi Dynamic Tiering](#) on the Hitachi Data Systems website.

Hitachi Compute Blade 2000

The Hitachi blade server model 2000 is an enterprise-class platform that offers the following features:

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

Hitachi embeds logical partitioning virtualization in the firmware of the Hitachi Compute Blade 2000 server blades. This proven, mainframe-class technology combines Hitachi's logical partitioning expertise with Intel VT technologies to improve performance, reliability, and security. Embedded logical partition virtualization does not degrade application performance and does not require the purchase and installation of additional components.

This solution takes advantages of logical partition features in the Hitachi Compute Blade 2000, such as N+1 cold stand-by design to establish a reliable and highly available virtualized Microsoft SQL Server solution. This solution scales from a 75,000 user deployment with a 900GB database to a total 150,000-user deployment by implementing a second 75,000 user deployment with a second 900GB database virtual machine and SQL server instance.

Microsoft SQL Server 2008 R2

The Microsoft website gives this overview for [Microsoft SQL Server 2008 R2](#):

Microsoft SQL Server 2008 R2 is a complete set of enterprise-ready technologies and tools that help people derive the most value from information at the lowest total-cost-of-ownership. Enjoy high levels of performance, availability, and security; employ more productive management and development tools; and deliver pervasive insight with self-service business intelligence (BI).

Microsoft SQL Server 2008 R2 also provides enhanced enterprise-class manageability for large OLTP deployments, like the one that is described in this reference architecture guide. When used with the Hitachi Virtual Storage Platform, Microsoft SQL Server 2008 R2 provides a scalable, high-performance database engine for any midrange to enterprise level application.

Hitachi Dynamic Link Manager Advanced

Hitachi Dynamic Link Manager Advanced is a software package that combines all the capabilities of Hitachi Dynamic Link Manager and Hitachi Global Link Manager into a comprehensive multipathing solution. Installed on each Microsoft SQL Server host, Dynamic Link Manager includes capabilities such as the following:

- Path failover and failback
- Automatic load balancing to provide higher data availability and accessibility.

The Hitachi Dynamic Link Manager Advanced, used for SAN multipathing, is configured using its round-robin load balancing algorithm. It automatically selects a path by rotating through all available paths. This balances the load across all available paths and optimizing IOPS and response time.

Tested Solution Components

The following provides details about the hardware and software components that were deployed in the Hitachi Data Systems lab.

Hardware Components

Table 1 lists detailed information about the hardware components.

Table 1. Hardware Components

<i>Hardware</i>	<i>Configuration</i>	<i>Version</i>	<i>Quantity</i>
Hitachi Virtual Storage Platform	<ul style="list-style-type: none"> ▪ 128GB cache memory ▪ 6 x 8Gb/sec front- end ports ▪ 35 x 300GB, 10K RPM, SAS disks ▪ 8 x 300GB, SSD disks ▪ 4 x 2TB, 7.2K RPM, SATA disks 	Microcode 70-00-63	1
Hitachi Compute Blade 2000 Chassis	<ul style="list-style-type: none"> ▪ 4 x 8Gb/sec dual-port Host Bus Adapters ▪ 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–X57A1 Blade	<ul style="list-style-type: none"> ▪ 2 x 8-Core Intel Xeon X7560 2.26GHz ▪ 128GB memory 	58.22	1
Brocade 5340 switch	SAN switch with 8Gb/sec Fibre Channel ports	FOS 6.4.0b	2

This solution used Hitachi Dynamic Tiering software to provision storage. Table 2 lists the storage configuration details for the dynamic tiering pools created for a building block supporting 75,000 users. RAID-5 (7D+1P) groups were used to create the dynamic tiering pools. This building block was replicated to support the 150,000 user environment.

Table 2. Dynamic Provisioning Pool for Single Virtual Machine Supporting 75,000 Users

<i>Pool ID</i>	<i>Number of Drives</i>	<i>Drive Type</i>	<i>Number of LUNs per Pool</i>	<i>Role</i>
01	16	300GB SAS 10K RPM	3	SQL Server database and tempdb files
	8	300GB SSD drives*		
02	8	300GB SAS 15K RPM	1	SQL Server transaction logs
03	4	2000GB SATA 7200 RPM	1	Logical partition OS in a shared pool with other logical partition OS LUNs

*A single LDEV comprising of a third of the total capacity of the SSD RAID group.

To scale this environment to support 150,000 users, add a second building block. This building block uses a 900GB database for each SQL server instance in a dynamically provisioned environment. The databases all use RAID-5 with a 7D+1P drive configuration for both the SAS and the SSD tier of the pool.

Table 3 lists the storage configuration details for logical divisions (LDEVs) created for the Microsoft SQL Server instance that supports a single SQL Server supporting 75,000 users. All LDEVs for the SQL Server instance were mapped to storage ports 1C, 3C, 5C, and 7C. The LDEV for the guest OS VHD was mapped to storage ports 3B and 4B. This 75,000 user building block was replicated to support a user environment supporting 150,000 total users.

Table 3. LDEV Storage Configuration for Single Virtual Machine Supporting 75,000 Users

<i>Pool ID</i>	<i>LDEV</i>	<i>LUN</i>	<i>Size (GB)</i>	<i>Description</i>
03	00:02:00	03	150	OS LUN for SQL Server logical partition operating system
01	00:20:10	04	450	Database LUN1 for SQL Server
01	00:20:11	05	450	Database LUN 2 for SQL Server
01	00:20:12	06	90	tempdb for SQL Server
02	00:20:02	07	225	Transaction logs for SQL Server

This LUN configuration is replicated on a second SQL Server guest machine to support the 150,000 user architecture.

Note— The LUN for the SQL server operating system is presented to the logical partition (LPAR) and used for storage area network boot.

Hitachi Compute Blade 2000 Blade Configuration

This solution uses the following:

- One Hitachi Compute Blade 2000 chassis
- One X55A2 blade
- Two Hitachi dual port 8Gb/sec host bus adapter (HBA) cards

The blade has two on-board NICs. Each NIC is connected to a LAN switch module.

The blade has two PCIe slots available. Each PCIe slot is populated with Hitachi dual port 8Gb/sec HBA cards. These cards are required for HVM mode. With Hitachi HBA cards, when HVM is enabled, eight virtual WWNs are created on each blade. Figure 2 shows the front and back view of Hitachi Compute Blade 2000 used in this solution.




Figure 2

To meet the 150,000 user requirement and to use the hardware resources efficiently, the Windows servers listed in Table 4 were deployed on one physical blade server.

Table 4. Hitachi Compute Blade 2000 Configuration

Blade Server	LPAR	Server Name	Role	Number of CPU Cores	Memory (GB)
Blade 1	LPAR1	SQLSERVER1	SQL Server instance	4	32
	LPAR2	SQLSERVER2	SQL Server instance	4	32
	LPAR3	AD2	Active Directory DNS	4	8

Each building block used to support 75,000 users is comprised only of a single LPAR. The active directory LPAR is optional.

Storage Area Network Components

To maintain high availability, 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.
- SQL Server volumes are connected to four HBA ports.
- Four redundant paths from the switches to the storage ports are configured for the SQL server volumes.
- Two redundant paths from the switches to the storage ports are used for SAN OS boot volumes.

The storage area network (SAN) configuration for this solution uses two Fibre Channel switches for high availability. Four redundant paths are configured from the SQL server host for the databases, `tempdb`, and transaction logs to the Hitachi Virtual Storage Platform.

For the SQL server logical partitions, the OS LUNs are configured with two redundant paths from the server blade to the Hitachi Virtual Storage Platform. The server blade has two dual port host bus adapters (HBAs) installed for high availability purposes.

Hitachi Dynamic Link Manager is used for multipathing, employing its round-robin multipathing policy. This load balancing algorithm automatically selects a path by rotating through all available paths. It balances the load across all available paths and optimizing IOPS and response time.

Figure 3 shows the Fibre Channel SAN architecture for the Microsoft SQL Server 2008 R2 implementation described in this solution.

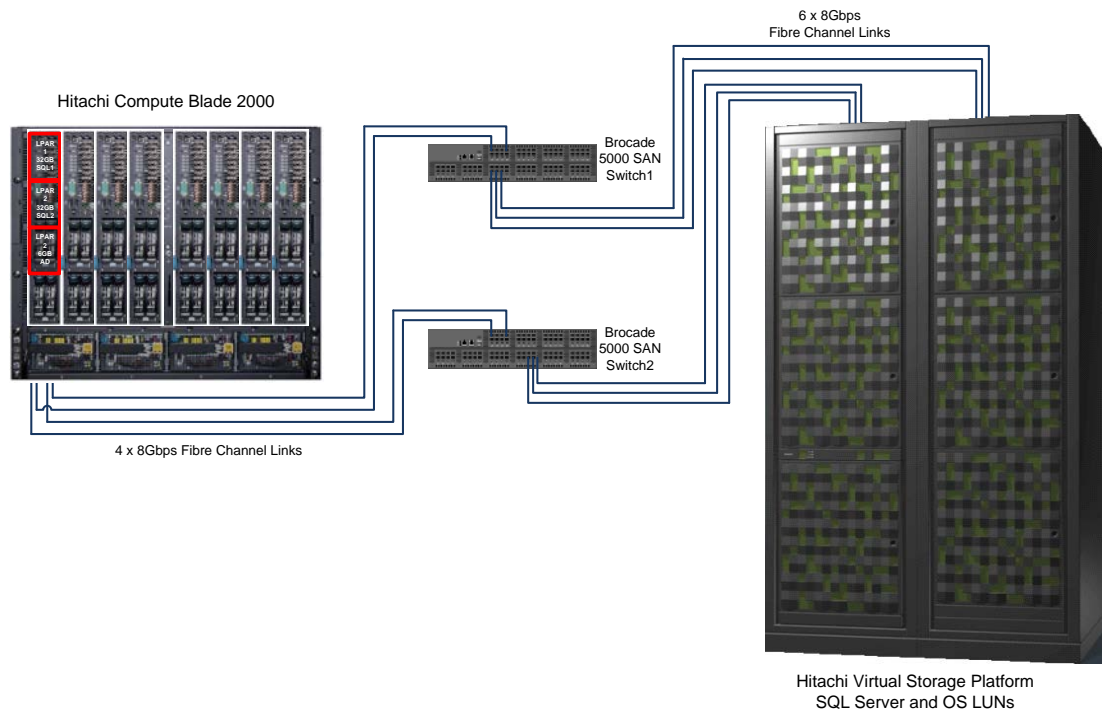


Figure 3

Software Components

Table 5 lists the information about the software components used in this solution.

Table 5. Software Components

<i>Software</i>	<i>Version</i>
Hitachi Storage Navigator	Dependent on microcode version
Hitachi Dynamic Tiering	Dependent on microcode version
Hitachi Dynamic Link Manager	6.5.0-00
Hitachi Global Link Manager	6.5.0-00
Microsoft Windows Server 2008 R2 (for SQL server and domain controller hosts)	Enterprise edition
Microsoft SQL Server 2008 R2	Enterprise edition

Solution Implementation

To deploy this Microsoft SQL Server 2008 R2 solution, follow these steps:

1. Configure the Hitachi Compute Blade 2000.
2. Configure fabric switch zones.
3. Configure the Hitachi Virtual Storage Platform.
4. Set up SAN OS boot.
5. Install Microsoft Windows 2008 on LPARs.
6. Install Hitachi Dynamic Link Manager.
7. Install Microsoft SQL Server 2008 R2.
8. Configure Microsoft SQL Server 2008 R2.

These general tasks need to be completed for a successful deployment. Your checklist can vary, based on your specific environment. The following sections have the details about each of these steps.

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

- **Management module web GUI**—This manages entire chassis using a web-based graphical user interface (GUI). To use it, open a browser and enter management module's IP address.
- **Blade server web GUI**—This manages a blade using a web-based GUI. To use it, open a browser and enter the blade's IP address.
- **Blade server command-line interface**—This manages HVM using a command-line interface. To use it, use a terminal emulator client to open a Telnet session to the blade's IP address.
- **Remote console**—This is the remote KVM for a blade. To use it, open a blade server web GUI. When the log on window opens, click the **Launch Remote Console** button.

Hitachi is changing the product names for its server products. Therefore, some of these interfaces may use old product names instead of the names used in this document.

For more information about each of these tasks, see the documentation provided by Hitachi Data Systems and Microsoft:

- Hitachi Compute Blade *User's Guide*
- Hitachi Virtual Storage Platform online help
- Hitachi Virtual Storage Platform Dynamic Tiering *User's Guide*
- Hitachi Dynamic Link Manager *User's Guide*
- Microsoft *TechNet* library for [SQL Server 2008 R2](#).

Implementation Preparations

Before implementing this solution, make sure you have the following items:

- 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)
- Seven to thirteen IP addresses:
 - Seven IP addresses
 - One for the server blade
 - Three for LPAR control (HVM)
 - Three for LPAR guest operating systems
 - Six IP addresses in environments where these items are not already in place and connected:
 - Two for the management module
 - Two for the storage systems
 - Two for the Fibre Channel switches

Make the following LAN connections:

- Connect one network cable from each network switch module to your corporate LAN, except for the first switch module. 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. SAN boot volumes for the OS are connected to two HBA ports. SQL Server volumes are connected to four HBA ports. Four redundant paths from the switches to the storage ports are configured for the SQL server volumes, and two redundant paths from the switches to the storage ports are used for SAN OS boot volumes.

Configure the Hitachi Compute Blade 2000

The following sections describe how to configure the Hitachi Compute Blade 2000. These procedures assume the following:

- The blade chassis is racked.
- The server blade is installed into the chassis.
- 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 laptop or personal computer for use as a system console 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. See Figure 4.




Figure 4

2. Configure your system console to use the following static IP address:
 - IP Address: 192. 168. 0. 100
 - Gateway: 192. 168. 0. 1
3. Open a browser and type the following in the address bar: `http://192. 168. 0. 1/`

A log on screen opens.

4. Log on using the default user name and password:
 - User name: **admini strator**
 - Password: **password**The management module web GUI opens.
5. Click the **Settings** tab.
6. In the navigation area on the left, click **The configuration of network link**.
The **Management LAN Network** area opens.
7. Click **Edit** in the **Management LAN network** area.
The fields become editable.
8. Enter an IP address, subnet mask, and default gateway for all of the following:
 - The management module
 - Each active partition (blade)You need one IP address for the management module and one IP address for each partition (blade).
9. Click **Confirm**, and then click **Apply**.
The IP settings are saved.
10. Disconnect the Ethernet cable from the management module.
11. Connect an Ethernet cable from the management module to an external management LAN switch. Then you can perform management through any system console in the management LAN.

Configure the Blade

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 type the blade IP address in the address bar.
2. Log on to the blade web GUI using the default user name and password:
 - User name: **user01**
 - Password: **pass01**The blade web GUI opens.
3. Click **EFI Setup**.
The **EFI Setup** area opens.

4. For the **SMT (Simultaneous Multi-Threading)** option, click **Disable**.
5. For the **PCI Error Handling Mode** option, click **Legacy**.
6. Keeping the remaining default values, click **Modify**.
The **EFI Setup (Confirm)** area opens.
7. Click **Confirm**.
The EFI settings are updated.
8. Click **HVM Setup**.
The **HVM Setup** area opens.
9. For the **OS Mode** option, click **HVM Mode**.
10. Click **Modify**.
A confirmation dialog box opens.
11. Click **Confirm**.
The dialog box closes and the **HVM Setup** area opens.
12. Click **Power and LEDs**.
The **Power and LEDs** area opens.
13. Click **Power on**.
The blade is turned on.

Configure HVM

Each partition requires HVM configuration. Follow these steps for each partition:

1. Initiate a Telnet session to the IP address assigned to the partition.
2. Log on using the default user name and password:
 - User name: `adminstrator`
 - Password: `password`.The **System Configuration** window opens with a message that indicates that HVM is being initialized. Wait until this message closes before continuing.
3. Place the cursor on **HVM IP Address** and press ENTER.
The **Change of HVM IP Address** menu opens.
4. Type the IP address of a partition and press ENTER.
5. Place the cursor on **Subnet Mask** and press ENTER.
The **Change of Subnet Mask** menu opens.
6. Enter the subnet mask and press ENTER.

7. Place the cursor on **Default Gateway** and press ENTER.

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

8. Enter 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 opens.

10. Type a different random number for each partition from 1 to 128 and press ENTER.

The VNIC system number creates a unique MAC address that does not conflict with other MAC addresses.

11. Press F10.

The confirmation menu opens.

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

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

13. Press ENTER.

The window closes.

14. Press ESC.

The main menu opens. Keep the Telnet session open for "Configure an LPAR."

Configure an LPAR

Configuring an LPAR involves assigning resources to it, such as CPU, memory, network, and Fibre Channel. This procedure uses the Telnet session started in "Configure." Repeat this procedure for each LPAR.

To set up an LPAR, follow these steps:

1. Place the cursor on **Logical Partition Configuration** and press ENTER.

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

2. Press F6.

The **Add LPAR** menu opens.

3. Place the cursor on the first **NO_NAME** and press ENTER.

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

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

The **Logical Partition Name** window opens.

5. Type the name for the logical partition and press ENTER.

For the first SQL server LPAR, Hitachi Data Solutions used *SQL-Server-01*.

- Place the cursor on the **Shr** column for shared CPU mode, and press ENTER.

The **Number of Shared Logical Processors** window opens.

- Type the number of shared CPUs for the LPAR and press ENTER.

For this implementation, each LPAR shared 4 CPUs. See Table 4.

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

The **memory size (in MB)** menu opens. See Figure 5.

```

+-----+
|+ 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 -----+| | | |
||                               | Within bounds | Per Memory : 4608 MB ||
|| Assign Total   5  0+-----+ Processors : 8(8)  ||
|| Act Total     0  0| Fl: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 |
+-----+

```

Figure 5

- Press F1.

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

- Type the memory size (in gigabytes) and press ENTER.

For the 150,000 user architecture, LPAR1 and LPAR2 were configured with 32GB. LPAR3 was configured with 8GB. See Table 4.

After pressing ENTER, the menu closes.

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

The **Auto Activation Order** menu opens.

Setting the auto activation order means each LPAR is automatically powered on in the specified order after you turn on the HVM.

- Type the order number for the LPAR and press ENTER.

The menu closes.

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

The **Virtual NIC Assignment** window opens.

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

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

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

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

25. Press ESC.

The main menu opens.

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

The **Shared FC Assignment** menu opens.

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

The **Shared FC Assignment** menu opens. See Figure 7.

```
+-----+
+- 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  ||                 1 ||
|| 4 LPAR4  Dea  ||                 2 ||
|| 5        ||                 3 ||
|| 6        ||                 4 ||
|| 7        ||                 5 ||
|| 8        ||                 6 ||
|| 9        ||                 7 ||
|| 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 ||
+-----+
```

Figure 7

28. Place the cursor on the first number available on the list and press ENTER.

Repeat Step 2728 and Step 28 for all shared Fibre Channel numbers.

29. Press ESC.

The main menu opens.

30. Press F9.

The message indicating that the configuration is being saved displays, possibly for several minutes. The configuration is saved when the **HVM configuration is saved successfully** message displays.

Hitachi Dynamic Link Manager is used for multipathing using 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 configuration used four Brocade 5300 SAN switches. If your environment uses different switches, some of the following steps might vary.

All zones configured on the switches using the single initiator approach.

- Configure two zones for each host for use of the SAN OS boot. Each zone must use a different switch to provide redundancy.
- Configure four zones for each host for use of the SQL server database, `tempdb`, and log. Each zone must use a different switch 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. Do one of the following:
 - Create a zone configuration and add the newly created zones to it.
 - Add the newly created zones to an already existent zone configuration.
6. Enable the zone configuration.

In addition, follow these best practices when configuring your SAN environment:

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

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

Table 6 lists the zoning deployed for the 75,000 user environment in the Hitachi Data Systems lab.

Table 6. 75,000 user environment

<i>LPAR Host</i>	<i>Switch</i>	<i>Host HBA Number</i>	<i>Zone Name</i>	<i>Storage System Port</i>	<i>Storage System Host Group</i>
LPAR 1	Brocade 5100-1	HBA 1 port 1	HCB2K_01_B0_L1_HBA1_1_ASE45_36_1C_4B	1C 4B	SQL_Server_1_HBA1_1
	Brocade 5100-2	HBA 1 port 2	HCB2K_01_B0_L1_HBA1_2_ASE45_36_7C	7C	SQL_Server_1_HBA1_2
	Brocade 5100-1	HBA 2 port 1	HCB2K_01_B0_L1_HBA2_1_ASE45_36_3C	3C	SQL_Server_1_HBA2_1
	Brocade 5100-2	HBA 2 port 2	HCB2K_01_B0_L1_HBA2_2_ASE45_36_5C_7B	5C 7B	SQL_Server_1_HBA2_2
LPAR 3	Brocade 5100-1	HBA 1 port 1	HCB2K_01_B0_L3_HBA1_1_ASE45_36_4B	4B	AD2_HBA1_1
	Brocade 5100-2	HBA 2 port 2	HCB2K_01_B0_L3_HBA2_2_ASE45_36_7B	7B	AD2_HBA2_2

Table 7 lists the zoning deployed for the 150,000 user environment in the Hitachi Data Systems lab.

Table 7. 150,000 user environment

<i>LPAR Host</i>	<i>Switch</i>	<i>Host HBA Number</i>	<i>Zone Name</i>	<i>Storage System Port</i>	<i>Storage System Host Group</i>
LPAR 1	Brocade 5100-1	HBA 1 port 1	HCB2K_01_B0_L1_HBA1_1_ASE45_36_1C_4B	1C 4B	SQL_Server_1_HBA1_1
	Brocade 5100-2	HBA 1 port 2	HCB2K_01_B0_L1_HBA1_2_ASE45_36_7C	7C	SQL_Server_1_HBA1_2
	Brocade 5100-1	HBA 2 port 1	HCB2K_01_B0_L1_HBA2_1_ASE45_36_3C	3C	SQL_Server_1_HBA2_1
	Brocade 5100-2	HBA 2 port 2	HCB2K_01_B0_L1_HBA2_2_ASE45_36_5C_7B	5C 7B	SQL_Server_1_HBA2_2
LPAR 2	Brocade 5100-1	HBA 1 port 1	HCB2K_01_B0_L2_HBA1_1_ASE45_36_1C_4B	1C 4B	SQL_Server_2_HBA1_1
	Brocade 5100-2	HBA 1 port 2	HCB2K_01_B0_L2_HBA1_2_ASE45_36_7C	7C	SQL_Server_2_HBA1_2
	Brocade 5100-1	HBA 2 port 1	HCB2K_01_B0_L2_HBA2_1_ASE45_36_3C	3C	SQL_Server_2_HBA2_1
	Brocade 5100-2	HBA 2 port 2	HCB2K_01_B0_L2_HBA2_2_ASE45_36_5C_7B	5C 7B	SQL_Server_2_HBA2_2
LPAR 3	Brocade 5100-1	HBA 1 port 1	HCB2K_01_B0_L3_HBA1_1_ASE45_36_4B	4B	AD2_HBA1_1
	Brocade 5100-2	HBA 2 port 2	HCB2K_01_B0_L3_HBA2_2_ASE45_36_7B	7B	AD2_HBA2_2

Configure the Virtual Storage Platform

The following describes how to configure the Hitachi Virtual Storage Platform for this solution. It assumes that the RAID groups and LDEVs to be used in the Dynamic Tiering Pool have already been setup.

Create Host Groups

Before you create a host group, you must enable port security. Port security is used to isolate traffic from multiple servers and group the traffic in common configurations.

Host group members must do the following:

- Run the same operating system
- Connect to the same port

You can do the following:

- Define up to 2,048 LUN paths for one host group
- Assign up to 255 host groups to one Fibre Channel port
- Assign up to 255 WWPNs to one Fibre Channel port

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

1. Open a browser and type the IP address of Storage Navigator in the address bar.
2. Log on using the assigned user name and password:
3. From the **Actions** menu, point to **Ports/Host Groups** and then click **Create Host Groups**.

The **Create Host Groups** window opens.

4. Type a name in **Host Group Name**.
5. From the **Host Mode** list, click **2C[Windows]**.
6. In the **Available Hosts** area, highlight one or more hosts.
7. In the **Available Ports** pane, select one or more ports.
8. Click **Add**.

The **Selected Host Groups** area is populated with the previously selected options.

9. Click **Finish**.

The **Create Host Groups** window opens.

10. Click **Apply**.

Creating a Dynamic Tiering Pool

You must create three pools on the Virtual Storage Platform, as follows:

- One pool for the operating systems
- One pool for the SQL Server databases and `tempdb`
- One pool for the SQL Server transaction logs

This procedure assumes that zoning within your Fibre Channel fabric is complete. Refer to Table 2 to establish the pools necessary to build this reference architecture.

Note —Although these steps mention the creation of Hitachi Dynamic Provisioning pool, they actually create the Hitachi Dynamic Tiering pool that is used for the reference architecture mentioned in this paper. The primary difference when creating a dynamic tiering pool is the enabling of the tiering capability along with the selection of LDEVs that use different drive types for the pool.

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

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

The **Create Pools** dialog box opens.

2. From the **Pool Type** list, click **Dynamic Provisioning**.
3. Click **Enable** for the **Multi-Tier Pool** option.
4. (Optional) Click an item from the **Drive Type/RPM** list and from the **RAID Level** list.

These options allow you to filter the available pool volumes.

5. Click **Select Pool VOLs**.

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

6. Select one or more pool volumes in the **Available Pool Volumes** area and click **OK**.
 - For the database and the `tempdb` pool, select two SAS LDEVs and one SSD LDEV. Make sure that they are configured as a RAID (7D+1P) group and that the SSD LDEV is only one-third of the overall RAID group capacity.
 - For the log pool, select a single LDEV comprised of a RAID (7D+1P) group with SAS drives.
 - For the SAN boot OS pool, select a single LDEV comprised of a RAID (3D+1P) group with SATA drives.

The **Create Pools** dialog box opens with **Total Selected Pool Volume** and **Total Selected Capacity** populated.

7. Type a prefix for the pool name in **Prefix**.

In the Hitachi Data Solutions environment, the following prefixes were used to identify how each pool is being utilized:

- Prefix for the OS Pool: HDT-MSFT-OS
- Prefix for Database Pool: HDT-MSFT-SQL-DB
- Prefix for Log Pool: HDT-MSFT-SQL-LOG

8. (Optional) Type an initial number for the pool name in **Initial Number**.

9. Click to expand the **Options** area.

10. Type a value in **Pool ID**.

11. Type a subscription limit in **Subscription Limit**.

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

12. Click a value from the **User-Defined Threshold** list and click **Add**.

The **User Defined Threshold** value determines when to trigger a pool capacity alert.

The **Selected Pools** pane is populated with the previously selected options.

13. Click **Finish**.

The **Create Pools** dialogue box opens.

14. (Optional) Click **Next** to start the creation of LDEVs within the pool as well selecting the host groups for assignment.

There is more detail on these operations in “Create” and “Mapping an LDEV.”

15. Click the **Apply** button.

Create LDEVs within a Dynamic Tiering Pool

This procedure assumes that dynamic pool creation in your environment is complete. Refer to Table 3 and Table 4 to establish the LDEVs necessary to build this environment.

To create one or more LDEVs within a Dynamic Tiering pool using Hitachi Storage Navigator software, follow these steps:

1. From the **Actions** menu, point to **Logical Device** and then click **Create LDEVs**.

The **Create LDEVs** window opens.

2. From the **Provisioning Type** list, click **Dynamic Provisioning**.

3. From the **Emulation Type** list, click **OPEN-V**.

4. For the **Multi-Tier Pool** option, click **Enable**.

5. (Optional) Choose an item from the **Drive Type/RPM** list and from the **RAID Level** list.

These options allow you to filter the available pools.

If using this option, only click an item from the **RAID Level** list for the DB and `tempdb` LDEVs. These are provisioned out of a pool with multiple tiers.

6. Click **Select Pool**.

The **Select Pool** window opens.

7. Click the pool in the **Available Pools** area and click **OK**.

- Click **HDT-MSFT-SQL-DB** when creating the database volumes (including `tempdb`) and the index volumes for the environment.
- Click **HDT-MSFT-SQL-LOG** when creating the log volumes for the environment.
- Click **HDT-MSFT-OS** when creating the SAN Boot OS volumes for the environment.

The **Create LDEVs** window opens with values for **Selected Pool Name** and **Selected Pool Capacity**.

8. Type the capacity amount in **LDEV Capacity** and click a unit of measure from the list.

9. Type the number of LDEVs of that size to be created in **Number of LDEVs**.

10. In the **LDEV Name** area, do the following:

- Type a prefix for the LDEV names in **Prefix**.
- Type an initial number for numbering the LDEVs in **Initial Number**.

11. Click to expand the **Options** area.

12. Modify the LDKC value if the default value of **00** is not appropriate.

You may need to change this if the storage is configured with more than one LDKC.

13. Click a value from the **CU** list.

14. Click a value from the **DEV** list.

15. (Optional) Click a value from the **Interval** list to not number the LDEVs sequentially.

Leave this value at the default of **0** to sequentially number the LDEVs.

16. Change the default values in **Initial SSID**, **CLPR**, and **Processor Blade** only if your environment requires different values.

17. Click **Add**.

The **Selected LDEVs** pane is populated with the previously selected options.

18. (Optional) Click **Next** to start the selection of the LDEVs for assignment to host groups.

If you choose this optional step, you will complete step 19 and step 20 of “Create LDEVs within a Dynamic Tiering Pool” after completing the full configuration of the environment in “Mapping an LDEV.”

19. Click **Finish**.

The **Create LDEVs** window opens.

20. Click **Apply**.

Mapping an LDEV

LDEV mapping makes LDEVs available to hosts and servers so that data can be stored and accessed by the host and its applications.

To map a LDEV using Hitachi Storage Navigator, 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**.

The **Selected LDEVs** pane is populated.

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**.

The **Selected Host Groups** pane is populated with the previously selected options.

7. Click **Next**.

The **Add LUN Paths** window opens.

8. Click **Finish**.

The **Add LUN Paths** window opens.

9. Click **Apply**.

Set Up SAN OS Boot

SAN boot is the capability to boot the OS from a LDEV on a storage system without using any of the server's internal disks. SAN boot is required for LPARs and the N+1 cold standby features on the Hitachi Compute Blade. The following sections describe how to set up the SAN OS boot.

Repeat these procedures for each LPAR.

Assign USB Devices and the Remote Console to an LPAR

Skip this section 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 for the blade that contains the LPARs.
2. On the main menu, place the cursor on **PCI Device Assignment** and press ENTER.

The **PCI Device Assignment** window opens.

3. Press F5.

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

4. Place the cursor on the LPAR that you want to attach the USB device to 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.

The PCI device is attached to the selected LPAR.

Activate an LPAR

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

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

1. On the main menu, press F3.

The **Activate LPAR** menu opens.

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 LPAR is powered on.

Leave the Telnet session open for “Configure HBAs.”

Configure HBAs

This procedure assumes that the Telnet session from “Activate an LPAR” is open.

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

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

A log on screen opens.

2. Click **Launch Remote Console**.

The remote console log on window appears.

3. Log on using the default user name and password:

▪ User name: **user01**

▪ Password: **pass01**

The remote console opens.

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

The EFI shell opens.

5. Type the following command at the prompt:

```
drivers
```

A list of drivers displays.

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

For the case in Figure 9, the driver number to record for later use 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 9

7. Type the following command at the prompt using the driver number recorded in step 6:

```
drvcfg <driver number>
```

For the example in Figure 9, you would type this: `drvcfg 70`.

A list of configurable components displays, as shown in Figure 10.

```
shell> 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 10

8. Record the driver number and controller number of the first configurable component on the list for later use.
9. Type the following command at the prompt using the driver number recorded in step 6 and the controller number recorded in step 8:

```
drvcfg -s <driver number> <controller number>
```

For the example in Figure 10, you would type this: `drvcfg -s 70 7B`

The prompt changes to **hfccfg>**.

10. Type the following command at the prompt:

```
select
```

A list of Fibre Channel ports and WWNs displays.

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

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

This solution uses Fibre Channel port 3.

12. Type the following command at the 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

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

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

y

Please select a number displays.

2) To select a WWN to configure, type the following at the prompt:

1

A list of options displays.

3) To scan for targets, type the following at the prompt:

1

A list of target devices displays.

4) To scan for LUNs, type the following at the prompt:

1

A list of LUNs displays.

5) To choose a LUN, type the following at the prompt:

1

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

Update list#1 (y/n) displays.

6) To update the list, type the following at the prompt:

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. Type the following command at the prompt:

save

A confirmation prompt displays.

16. To confirm the save command, type the following at the prompt:

`y`

A confirmation message displays. The path to the storage for the boot OS is configured.

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

17. To conclude, type the following command at the prompt:

`exit`

You are returned to the main EFI shell. Leave the remote console session open for “Map OS Boot Device.”

Map OS Boot Device

These instructions assume that the remote console session from “Configure HBAs” is open.

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

1. Type the following command at the prompt:

`reconnect -r`

A success message displays.

2. To remap the devices, type the following command at the prompt:

`map -r`

A device mapping table displays.

3. Type the following command at the prompt:

`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 opens.

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

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

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

The **File Explorer** window opens. See Figure 11.

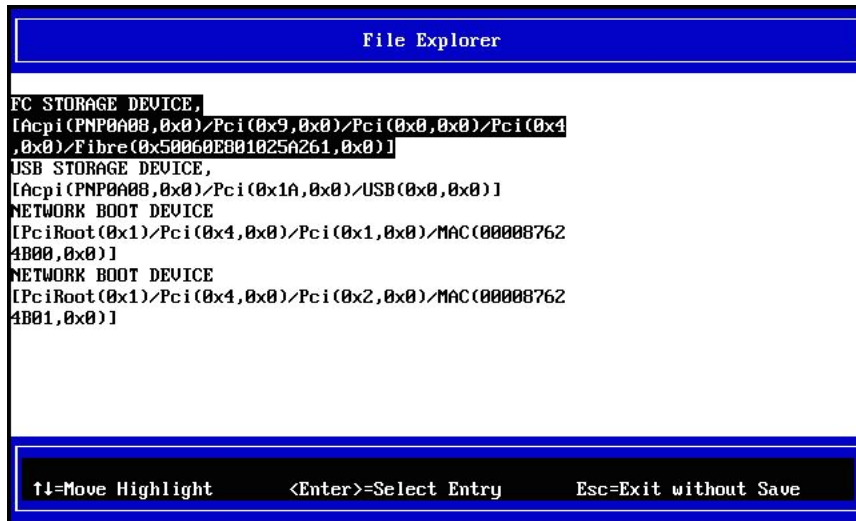


Figure 11

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

Multiple input display.

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

Input the description becomes editable. See Figure 12

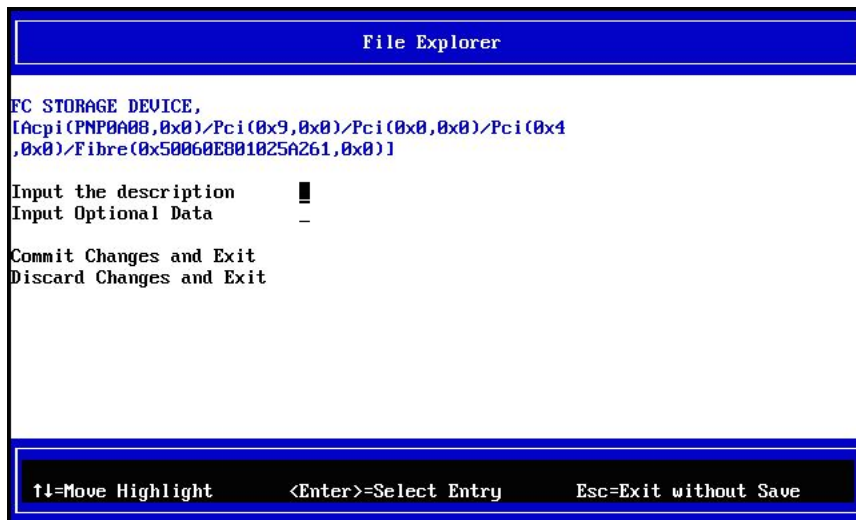


Figure 12

- Type a name for the device and press ENTER.

The prompt closes.

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

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

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

This procedure assumes that the remote console session from the “Map OS Boot Device” procedure is open.

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

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

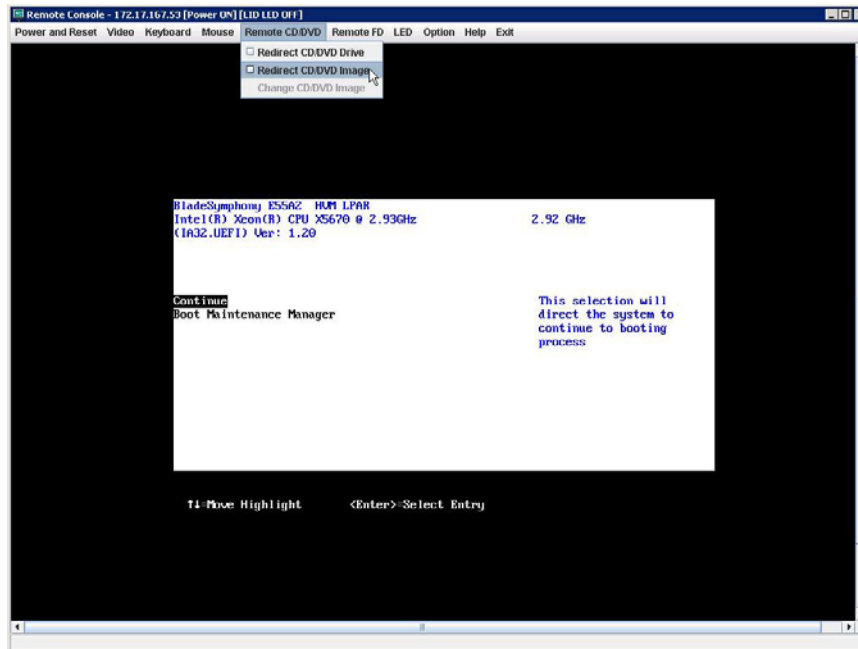


Figure 13

The **Open** window opens.

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 opens.

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

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

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

The **File Explorer** window opens.

6. Verify 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:
 - 1) Return to the server blade Command Prompt window.
 - 2) To select a LPAR to be reactivated, on the main menu, press F5.
Soon after the reactivation, the EFI main screen opens on remote console.
 - 3) Return to step 3 and repeat.
7. Place the cursor on **USB STORAGE DEVICE** and press ENTER.
Multiple inputs 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 values in the **Boot Maintenance Manager** window change.
12. Place the cursor on **Change the order** and press ENTER.
The **Boot Order** menu opens.
13. Use PLUS SIGN (+) and MINUS SIGN (–) to change the boot order to the following:
 - 1) Remote DVD
 - 2) The Fibre Channel storage LUNWhen finished, press ENTER.
14. Place the cursor on **Commit Change and Exit** and press ENTER.
This saves your changes and returns you 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 LPAR begins to boot from the Microsoft Windows ISO image.

Install Windows 2008 R2 on an LPAR

After the LPAR boots from the ISO image in “Configure LPAR to Boot Windows from an ISO Image on a Remote DVD,” follow the instructions from Microsoft to complete the installation, including:

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

You must install Microsoft Windows Server 2008 R2 on all active LPARs. Verify that all Windows updates are applied on each installation.

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

Install Active Directory

For this solution, one domain controller was installed on the LPAR 2 chassis of the blade being used. This optional installation was done to provide active directory services to the SQL servers that are part of this architecture. When installing in a production environment, add the domain controller to the existing Active Directory infrastructure.

To install Active Directory on a server, follow the instructions from Microsoft. 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

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

Install Microsoft SQL Server 2008 R2

To install Microsoft SQL Server 2008 R2, follow the instructions from Microsoft in *TechNet* [SQL Server 2008 Library](#) under the SQL Server 2008 R2 section.

Configure Microsoft SQL Server 2008 R2

The following describe storage configuration requirements for Microsoft SQL Server when hosting the databases, logs, and `tempdb` files on the Hitachi Virtual Storage Platform.

Databases and Transaction Log Files

When provisioning the storage for the SQL databases and logs, establish an allocation strategy using these configuration parameters:

- **SIZE**
- **FILEGROWTH**
- **MAXSIZE**

Use the parameters illustrated in the following SQL code sample when creating the databases and transaction logs with Hitachi Dynamic Provisioning, changing the placeholders shown in red to the values for your environment:

```
CREATE DATABASE <name of database> ON PRIMARY (  
    NAME = <logical database name>,  
    FILENAME = <OS location and name of database>,  
    SIZE = <initial size of the database file in MB>,  
    MAXSIZE = <maximum size to which the database file can grow in MB>,  
    FILEGROWTH = <specifies the database growth increment in MB>)  
LOG ON (  
    NAME = <logical log file name>,  
    FILENAME = <OS location and name of log file>,  
    SIZE = <size for log file in MB>,  
    MAXSIZE = <maximum size to which the log file can grow in MB>,  
    FILEGROWTH <specifies the log growth increment in MB>)
```

For this implementation, the Hitachi Data Systems lab used the following SQL statements, with lab-specific parameters shown in red:

```
CREATE DATABASE mydb ON PRIMARY (  
    NAME = mydb,  
    FILENAME = 'E:\mountpoint\data1\mydb1.mdf',  
    SIZE = 420MB,  
    MAXSIZE = 42000,  
    FILEGROWTH = 42)  
LOG ON (  
    NAME = mydblog1,  
    FILENAME = 'L:\mountpoint\log\mydblog1.ldf',  
    SIZE = 420MB,  
    MAXSIZE = 10000,  
    FILEGROWTH = 42)
```

Each file is created with an initial **SIZE** and auto extends by **FILEGROWTH** when the currently allocated space runs out. It stops the auto-extend if it reaches **MAXSIZE**.

For the database, the values of **SIZE**, **FILEGROWTH** and **MAXSIZE** are determined by expected growth, as follows:

- Set **SIZE** to a value that creates an initial allocation of data that is at least equal to one of the following:
 - The amount of data you plan to import into the database
 - The initial amount of data you expect for the database or log files
- Verify that the value for **FILEGROWTH** is a multiple of the size of a page, 42MB, which is the ideal value when using Hitachi Dynamic Provisioning.
- **MAXSIZE**, a more straightforward parameter, is used to keep your application from exceeding capacity, fail, or both.

Pay special attention to the **SIZE** and **FILEGROWTH** values given to log files to avoid the following for a SQL server instance:

- A **SIZE** that is too small. This might affect system performance.
- Repeated small increments from **FILEGROWTH** may allow the log files to grow too large. This can slow database startup, as well as log backup and restore operations.

Microsoft recommends that you assign the following to a log file:

- A value for **SIZE** this is close to the final size required.
- A relatively large value for **FILEGROWTH**. (Use a multiple of 42MB for implementations that use Hitachi Dynamic Tiering or Hitachi Dynamic Provisioning.)

Use the instant file initialization feature of SQL server to allow the following:

- Faster and optimized data file creation and growth
- Faster execution of database or filegroup restore operations

Instant file initialization reclaims used disk space without filling that space with zeros. Instead, disk content is overwritten as new data is written to the files. This makes it an ideal SQL server feature to work in conjunction with Hitachi Dynamic Provisioning. For more information, see the [Data File Initialization](#) page of Microsoft SQL Server 2008 *Books Online*.

tempdb Files

The default location for *tempdb* files in SQL server is on the drive C. You can modify the existing *tempdb* file location to point to the LUN that is provisioned on the Hitachi Virtual Storage Platform.

Modify the *tempdb* file location by executing the following SQL statements, changing the placeholders shown in red to the values for your environment:

```
USE tempdb
GO
ALTER DATABASE tempdb
MODIFY FILE (NAME='tempdev', FILENAME= <OS location and name of tempdb>,
SIZE = <size of file in KB or MB>,
FILEGROWTH = <specifies the tempdb growth increment in MB>)
```

For this implementation, the Hitachi Data Systems lab used the following SQL statements, with lab-specific parameters shown in red:

```
USE tempdb
GO
ALTER DATABASE tempdb
MODIFY FILE (NAME='tempdev', FILENAME= 'T:\SQL\DATA\datatempdb1.mdf',
SIZE=8192 KB, FILEGROWTH=42MB)
```

This locates the files on a dedicated LUN that is assigned to the SQL server for each tempdb.

 **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-092-00 July 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