HITACHI
DATA SYSTEMS

Hitachi TagmaStore®
Adaptable Modular Storage and Workgroup Modular Storage

VMware™ Configuration Guide

MK95DF742-02

Trademarks

Notice of Export Controls

Export of technical data contained in this document may require an export license from the United States government and/or the government of Japan. This technical data may not be used directly or indirectly (without U.S. Government permission), in the design, development, fabrication, production, testing or use of rocket systems, or unmanned air vehicles; or chemical or biological weapons; or nuclear weapons or explosives or operation of facilities or components thereof for chemical processing of irradiated special nuclear material, heavy water production, separation of isotopes of source and special nuclear material, or fabrication of nuclear reactor fuel containing plutonium, or in International Atomic Energy Agency (IAEA) un-safeguarded nuclear facilities.

Please contact the Hitachi Data Systems Legal Department for any export compliance questions.

Document Revision Level

Revision	Date	Description
MK-95DF742-00	June 2005	Initial Release
MK-95DF742-01	August 2005	Revision 01, supersedes and replaces MK-95DF742-00
MK-95DF742-02	June 2006	Revision 02, supersedes and replaces MK-95DF742-01

Changes in this Revision

- Added Workgroup Modular Storage to the title page.
- Changed the footer to read "Hitachi TagmaStore® AMS and WMS VMware™ Configuration Guide."

Preface

This VMware™ configuration guide describes how to install and configure devices on the Hitachi TagmaStore® Adaptable Modular Storage and Workgroup Modular Storage systems system for operation with a VMware ESX Server™ operating system. This document is intended for system administrators and Hitachi Data Systems representatives. If you have questions or concerns, please contact the Hitachi Data Systems Support Center.

This document assumes the following:

- The user has a background in data processing and understands disk storage systems and their basic functions, and
- The user is familiar with the Hitachi TagmaStore Adaptable Modular Storage or Workgroup Modular Storage system, and
- The user is familiar with the VMware operating system, the VMware server, and the fibre-channel adapters.

For more information about the Hitachi Adaptable Modular Storage solutions, either contact Hitachi Data Systems or refer to the following documents:

- Adaptable Modular Storage
  - Adaptable Modular Storage 200 system
    Hitachi TagmaStore Adaptable Modular Storage 200 User and Reference Guide, MK-95DF713
  - Adaptable Modular Storage 500 system
    Hitachi TagmaStore Adaptable Modular Storage 500 User and Reference Guide, MK-95DF714
  - Adaptable Modular Storage 1000 system
    Hitachi TagmaStore Adaptable Modular Storage 1000 User and Reference Guide, MK-95DF780
- Workgroup Modular Storage
  - Workgroup Modular Storage 100 system
    Hitachi TagmaStore Workgroup Modular Storage 100 User and Reference Guide, MK-95DF738For more information about the VMware operating system, consult the VMware user documentation, or contact VMware technical support.

Note: The use of the Hitachi TagmaStore Adaptable Modular Storage and all Hitachi Data Systems products is governed by the terms of your agreement(s) with Hitachi Data Systems.

Microcode Level

This document revision applies to TagmaStore Adaptable Modular Storage and Workgroup Modular Storage versions 2.0 and higher.

Convention for Storage Capacity Values

Storage capacity values for hard disk drives (HDDs) on the AMS and WMS systems are calculated based on the following values:

1 KB = 1,000 bytes
1 MB = $1,000^2$ bytes
1 GB = $1,000^3$ bytes
1 TB = $1,000^4$ bytes

Storage capacity values for logical units (LUs) on the AMS and WMS systems are calculated based on the following values:

1 KB = 1,024 bytes
1 MB = $1,024^2$ bytes
1 GB = $1,024^3$ bytes
1 TB = $1,024^4$ bytes

Referenced Documents

Hitachi Data Systems documentation:

- *Hitachi TagmaStore Adaptable Modular Storage 200 User and Reference Guide*, MK-95DF713
- *Hitachi TagmaStore Adaptable Modular Storage 500 User and Reference Guide*, MK-95DF714
- *Hitachi TagmaStore Adaptable Modular Storage 1000 User and Reference Guide*, MK-95DF780
- *Hitachi TagmaStore Workgroup Modular Storage 100 User and Reference Guide*, MK-95DF738
- *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage: Storage Navigator Modular Graphical User Interface (GUI) User's Guide*, MK-95DF711
- *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage: Storage Navigator Modular Command Line Interface (CLI) User's Guide*, MK-95DF712

VMware documentation:

- VMware ESX Server Installation Guide: http://www.VMware.com/support/pubs/esx_pubs.html
- VMware ESX Server Administration Guide: http://www.VMware.com/support/pubs/esx_pubs.html

Comments

Please send us your comments on this document. Make sure to include the document title, number, and revision. Please refer to specific section(s) and paragraph(s) whenever possible.

- **E-mail:** doc.comments@hds.com
- **Fax:** 858-695-1186
- **Mail:**
  Technical Writing, M/S 35-10
  Hitachi Data Systems
  10277 Scripps Ranch Blvd.
  San Diego, CA 92131

Thank you! (All comments become the property of Hitachi Data Systems Corporation.)

Contents

List of Figures

List of Tables

Chapter 1 Introduction

Welcome to the VMware™ Configuration Guide for Hitachi TagmaStore® Adaptable Modular Storage and Workgroup Modular Storage (hereinafter referred to as AMS/WMS).

This chapter covers the following topics:

- Configuration Road Map (section 1.1)
- Adaptable Modular Storage/Workgroup Modular Storage System (section 1.2)

This document describes how to:

- Connect the Hitachi TagmaStore AMS/WMS to a VMware ESX Server.
- Configure the new Hitachi TagmaStore AMS/WMS devices for operation with the VMware operating system.

Some of these activities are performed by the Hitachi Data Systems representative, while other activities are performed by the user.

- The Hitachi Data Systems representative performs the physical installation of the Hitachi TagmaStore AMS/WMS.
- The user prepares for the installation and configures the new TagmaStore AMS/WMS devices, with assistance from the Hitachi Data Systems representative.

1.1 Configuration Road Map

Configuring the TagmaStore AMS/WMS disk devices for VMware operations involves the following tasks:

- Verifying VMware new storage device (section 4.1)
- Creating VMware VMFS-2 Volume (section 4.2)
- Assigning VMFS-2 and raw devices to Guest OS (section 4.3)
- Managing Guest OS (section 4.4)
- Verifying Guest OS access (section 4.5)
- Online LUN Addition and Deletion (section 4.6)

Note on the term "SCSI disk": The AMS/WMS logical devices are defined to the host as SCSI disk devices, even though the interface is fibre-channel.

1.2 Adaptable Modular Storage and Workgroup Modular Storage System

The Hitachi AMS/WMS systems are high-performance storage arrays, with added features for increasing data accessibility and enabling continuous user data access.

- The AMS systems offer a mix of both fibre-channel and lower-cost Serial ATA (SATA) drives.
- The WMS systems can be configured with SATA drives only for low-cost near-line storage applications, which could be used for purposes such as archiving data for regulatory compliance.

The AMS/WMS architecture empowers users to scale the system to meet a wide range of capacity and performance requirements. In addition, the AMS/WMS provides connectivity to most open systems through a standard fibre-channel interface.

For more information about the AMS/WMS systems, contact Hitachi Data Systems or refer to the following documents:

- AMS Documents
  - *Hitachi TagmaStore Adaptable Modular Storage 200 User and Reference Guide*, MK 95DF713
  - *Hitachi TagmaStore Adaptable Modular Storage 500 User and Reference Guide*, MK 95DF714
  - *Hitachi TagmaStore Adaptable Modular Storage 1000 User and Reference Guide*, MK-95DF780
- WMS Document
  - *Hitachi TagmaStore Workgroup Modular Storage 100 User and Reference Guide, MK-95DF738*

Notes:

- The availability of the AMS/WMS features and functions depends on the level of microcode installed.
- Logical devices are defined to the host as SCSI disk devices, even though the interface is fibre-channel.

Chapter 2 Preparing for New Device Configuration

This chapter describes how to prepare for new device configuration. Topics include:

- Configuration requirements (section 2.1)
- Installing the TagmaStore Adaptable Modular Storage/Workgroup Modular Storage (section 2.2)
- Host Mode options (section 2.3)
- Preparing to connect to the TagmaStore Adaptable Modular Storage/Workgroup Modular Storage (section 2.4)
- Connecting the TagmaStore Adaptable Modular Storage/Workgroup Modular Storage to the VMware server (section 2.5)
- Configuring the host fibre-channel adapters (section 2.6)

2.1 Configuration Requirements

This section lists the minimum hardware, system, and software requirements for the SAN configuration. For a complete list of requirements, refer to the VMware ESX Server™ Installation Guide or visit http://www.VMware.com/support.

- **VMware Intel® Base ESX Server**
  - Two to 16 processors:
    - 900MHZ Pentium® III Xeon and above
  - 512 MB of system memory
  - Two supported Ethernet controllers:
    - Broadcom NetXtreme 570x gigabit controller
    - Intel® PRO/100 and 1000 adapters
    - 3Com 9xx-based adapters
  - One internal SCSI, fibre-channel, or internal RAID controller with one internal hard drive:
    - Adaptec® and LSI SCSI, IBM®, HP, and Dell™ RAID controllers
    - Emulex® or QLogic® fibre-channel controllers
  - Two fibre-channel host bus adapters (HBAs):
    - QLogic HBA 2342, 2310 with firmware v1.34
    - Emulex LP9002 and LP9802 HBA
- **VMware ESX Server OS and Management Interface**
  - VMware ESX Server software version 2.5, build 11343
  - Standard x86-based computer with Web browser capability (e.g., Internet Explorer 6)
  - VMware Management Interface accessed from ESX Server
  - Guest OS for VMware virtual machines. Hitachi supports:
    - Guest OS Windows® 2000 Service Pack 4
    - Windows® 2003 EE
    - Linux® Red Hat® AS 3.0 UPD3
    - Linux Red Hat AS 2.1 UPD5

    Note: For the latest VMware Guest OS support, refer to the HDS support document or contact Hitachi Data Systems.
  - License numbers for ESX Server and VMware Virtual Machine SMP for ESX Server

- **Hitachi Storage System**
  - The AMS/WMS software is required to configure the fibre-channel ports. If the remote Storage Navigator Modular feature is not installed, contact Hitachi Data Systems for information about LUN and fibre-channel configuration services.
  - The AMS/WMS software is required to configure the fibre-channel ports.
  - The AMS/WMS supports 1, 2, and 4 Gigabit fibre-channel interfaces.

    Note: Do not connect any OFC-type fibre-channel interface to the AMS/WMS. For information about supported fibre-channel (FC) adapters, optical cables, hubs, and fabric switches, contact Hitachi Data Systems.
- **Storage Area Network (SAN)**

  The AMS/WMS is required to connect to VMware ESX Server via a SAN. For information about supported fibre-channel switches, topologies, and firmware versions, contact Hitachi Data Systems.

  Note: VMware does not support direct connection to storage system.
- **Local Area Network**

  A 100BaseT or faster network-interface card.

Note: The VMware administrator password is required for TagmaStore AMS/WMS device configuration.

2.2 Installing the Adaptable Modular Storage System

The AMS/WMS comes with all the hardware and cabling required for installation. Installation of the AMS/WMS involves the activities in the following sections.

2.2.1 Hardware Installation

Perform the hardware installation as specified in the Hitachi Data Systems documentation. The hardware-installation procedure includes:

- Following all precautions and procedures in the Hitachi documentation,
- Checking all specifications to ensure proper installation and configuration,
- Assembling all hardware and cabling,
- Upgrading to the latest microcode level,
- Creating RAID groups and LUNs and formatting LUNs using the Storage Navigator Modular Storage software. For more information, refer to the following documents:
  - *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage: Storage Navigator-Modular Command Line Interface (CLI) User's Guide* (MK-95DF712)
  - *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage: Storage Navigator-Modular Graphical User Interface (GUI) User's Guide* (MK-95DF711)
  - *Hitachi TagmaStore Adaptable Modular Storage Model AMS500 User and Reference Guide* (MK-95DF714)
- Installing the fibre-channel adapters and cabling,

2.2.2 TagmaStore Adaptable Modular Storage FC Port

The fibre topology parameters for each AMS/WMS fibre-channel port depend on the type of device to which the AMS/WMS port is connected. Determine the topology parameters supported by the device, and set your topology accordingly (refer to section 2.4.2).

2.3 Host Mode Options

To provide flexibility and enable the AMS/WMS to be tailored to unique customer operating requirements, additional operating parameters called “Host Mode options” are available for the AMS/WMS.

At installation, the Host Mode options are set to default values, which may or may not be appropriate for your requirements. It is, therefore, important to contact Hitachi Data Systems about settings that are applicable to the VMware environment and best suited for your applications. Only Hitachi Data Systems representatives can change Host Mode options.

2.4 Preparing to Connect TagmaStore Adaptable Modular Storage/Workgroup Modular Storage

Before the AMS/WMS is connected to your VMware server:

1. Set the host mode for the AMS/WMS fibre-channel port(s) (see section 2.4.1).
2. Configure the AMS/WMS fibre-channel ports (see section 2.4.2).
3. Use the Storage Manager Modular software to set the host specific parameters for the AMS/WMS and configure the fibre ports. For more information, refer to the *Hitachi TagmaStore Adaptable Modular Storage and Workgroup Modular Storage Navigator Modular User's Guide* (MK-94RD203).

After completing these steps:

1. Shut down the VMware system, connect the AMS/WMS, and restart the VMware system (see section 2.5).
2. Configure the FC adapters connected to the AMS/WMS (see section 2.6).

Note: If the remote Storage Navigator Modular feature is not installed, contact Hitachi Data Systems for information about fibre-channel configuration services.

2.4.1 Setting the Host Connection Mode for Ports

The AMS/WMS ports have host connection modes that must be set for the connected operating system. Use the System viewer software to set the host mode for each port.

The required host connection mode settings for VMware are:

- Standard Mode for hostmode 1.
- <disable> for hostmode 2.

In addition, the fibre channel topology must be set to point-to-point.

The following example describes how to use the System viewer file to change from Normal Mode to Management Mode:

1. Click **Logical Status**.
2. Select the host group.
3. Select the controller and the port to which your server is connected.
4. Select the configuration settings (in this example, Port 0A).
5. Select the host group.
6. Right click the option.
7. Click the **Modify Host Group** option.
8. Click **Details**.
9. Select **Standard** mode for host connection mode 1 and <disable> for hostmode 2.
10. Click **OK**.

2.4.2 Fibre Topology

The AMS/WMS ports must be configured for the connected operating system.

Configure the AMS/WMS fibre-channel ports to define the fibre topology parameters. The AMS/WMS supports up to 512 LUs and 128 host groups per port, with 256 LUs mapping per host group. You select the appropriate settings for each AMS/WMS FC port based on the device to which the port is connected. Determine the topology parameters supported by the device, and set your topology accordingly.

Notes:

- If you plan to connect different types of servers to the AMS/WMS via the same fabric switch, you must use either **zoning** on the switch or the Host Storage Domain on the AMS/WMS, or a combination of both.
- Contact Hitachi Data Systems for information about port topology configurations supported by each HBA/switch combination.

Figure 2.1 shows the **Storage Navigator Modular** window used to define the port parameters, and Table 2.1 describes the settings on this window. You select the appropriate settings for each FC port based on the device to which the port is connected. Determine the topology parameters supported by the device and set your topology accordingly. The type of port is also important.

Note: Although the AMS/WMS supports both fabric point-to-point and Arbitrated Loop topologies, VMware supports only a fabric point-to-point topology. FC-AL and direct-connect topologies are not supported.

Table 2.1 Fibre Parameter Settings for Universal Storage Platform

Fabric	Connection Parameter	Supported Not Supported
Enable	FC-AL	Not supported
Enable	Point-to-Point	Supported

Note: Contact Hitachi Data Systems for information about port topology configurations supported by each HBA/switch combination.

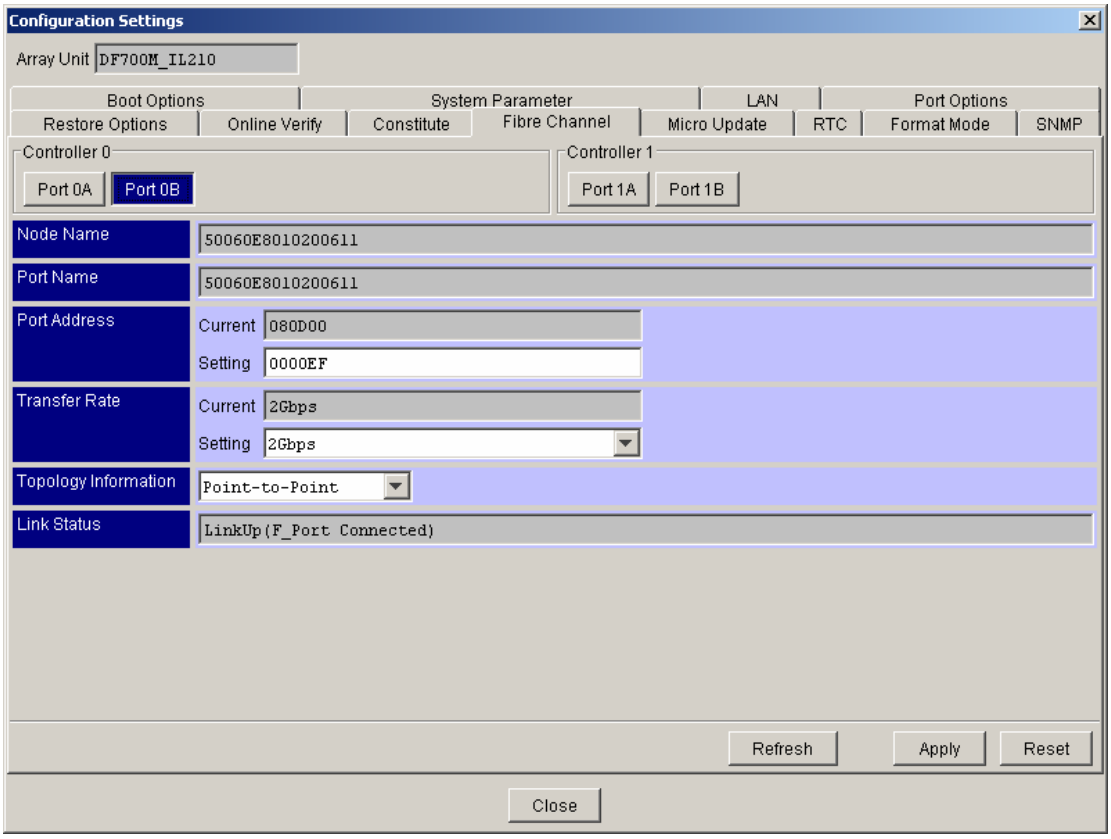

Figure 2.1 Verifying FC Port Topology

2.5 Connecting TagmaStore Adaptable Modular Storage to VMware Server

After you configure the AMS/WMS fibre-channel ports, connect the AMS/WMS to the VMware system. Some of these steps are performed by the Hitachi Data Systems representative, while other steps are performed by the user.

Note: The Hitachi Data Systems representative must use the Adaptable Modular Storage maintenance manual during all installation activities. Follow all precautions and procedures in the maintenance manual, and always check all specifications to ensure proper installation and configuration.

To connect the AMS/WMS:

1. **Verify system installation**. The Hitachi Data Systems representative verifies the fibre-port address configuration, and verifies that the status of the fibre-channel adapters and logical device is NORMAL.
2. **Shut down the VMware system**. The user shuts down and powers off the VMware system before connecting the AMS/WMS:
   a) If the VMware ESX Server is not at power-off state, shut down the VMware system by executing the following command at the service console: `Shutdown –h now`
   b) When shutdown is complete, power off the VMware display.
   c) Power off all peripheral devices, except the AMS/WMS.
   d) Power off the VMware system. You can now connect the AMS/WMS.
3. **Connect the AMS/WMS to the VMware system**. The Hitachi Data Systems representative installs the fibre-channel cables between the AMS/WMS and the VMware system.
4. **Power on and Boot the VMware system**. The user powers on the VMware system after connecting the AMS/WMS:
   a) Power on the VMware system display.
   b) Power on all peripheral devices. The AMS/WMS should already be on and the fibre-channel ports should already be configured. If the fibre-channel ports are configured after the VMware system is powered on, the system may need to be restarted for the new device to be recognized.
   c) Confirm the ready status of all peripheral devices, including the AMS/WMS.
   d) Power on the VMware system connected to the AMS/WMS.

2.6 Configuring the Host Fibre-channel Adapters

After connecting the AMS/WMS and rebooting the VMware server, configure the fibre-channel adapter(s) connected to the AMS/WMS. For information about configuring your fibre-channel adapter(s), refer to the documentation that came with your adapter(s).

The user documentation for your fibre-channel adapter(s) should also describe whether other options are required to meet your operational requirements. In particular:

- Check whether the BIOS needs to be disabled to prevent the system from booting from the Universal Storage Platform.
- Set the queue depth, BIOS, Target reset, and fibre-channel topology parameters.

Note: Use the same settings and device parameters for all AMS/WMS devices.

After configuring the fibre-channel adapter(s), you may have to reset VMware ESX Server to have the change take effect before booting up to VMware ESX Server.

2.6.1 ESX Server Guidelines

VMware ESX Server can be used with SANs. The ESX Server supports QLogic and Emulex® HBAs. These adapters allow an ESX Server computer to connect to a SAN and recognize the disk arrays on the SAN.

The SCSI configuration information contained in this section also applies to fibre-channel adapters, however, fibre-channel adapters may require additional configuration. For information about supported SAN hardware, download the VMware ESX Server SAN Compatibility List from the VMware Web site: *http://www.VMware.com/support*/resources/esx_resources.html.

2.6.2 Configuring the QLogic HBA BIOS

The default ESX Server driver is based on the QLogic-provided qla2x00 v6.07 driver. This driver has been tested with the ESX Server in various SAN fabric configurations. FC-AL configurations are currently not supported. The firmware for these devices is loaded by the driver. Point-to-point or direct connections are not supported unless noted otherwise.

Hitachi supports the following Emulex HBAs with VMware ESX Server 2.5:

- QLA-2310/2312 (Blade servers)
- QLA-2340/2342

For more information about supported QLA HBA models, contact Hitachi Data Systems. A list of VMware supported HBA can be found at http://www.VMware.com/pdf/esx_SAN_guide.pdf

To configure the QLogic HBA with VMware ESX Servers for Hitachi Storage solutions, use the following procedure. Other QLogic parameters may be configured depending on the specific environment:

1. Reboot the server.
2. During system startup, press **CTL+Q** to access the QLogic adapter settings menu.
3. Enter the QLogic HBA configuration utility during boot up.
4. Access the **Select Host Adapter** page, select the HBA to be configured, and press **Enter.**
5. From the **Configuration Settings**, enter **Restore Default Settings**.
6. From **Host Adapter Settings**, verify the following settings:
   - Bios Revision: 1.34 as the minimum level (check the VMware documentation for current QLogic firmware support).
   - Host Adapter BIOS: Disabled
   - Connection Options: 1 (Point-to-Point only) Default value is 2, press **Enter** and change the value to 1.
7. Return to the **Configuration Settings**, access the **Advanced Adapter Settings** window and verify the following settings:
   - LUNs per Target: a number that ranges with the number of LUNs of your configuration.

     Note: This setting specifies the number of LUNs per target. Multiple LUN support is typically for redundant array of independent disks (RAID) boxes that use LUNs to map drives. The default is 8. If you do not need multiple LUN support, set the number of LUNs to 0.
   - Enable LIP Reset: **No**
   - Enable LIP Full Login: **Yes**
   - Enable Target Reset: **Yes**
8. Save your changes and repeat the configuration for other HBA.

2.6.3 Configuring Emulex HBA BIOS

The ESX Server drivers are based on the Emulex-provided lpfcdd-2xx 2.01g driver. This driver has been tested with the ESX Server in various fabric configurations. FC-AL configurations are not supported. Point-to-point and direct connections are not supported unless noted otherwise.

Hitachi supports the following Emulex HBAs with VMware ESX Server 2.5:

- LP9002L (Minimum firmware revision 3.90a7)
- LP9802DC (Minimum firmware revision 1.00a4)

For more information, contact Hitachi Data Systems. A list of VMware-supported HBAs can be found at http://www.VMware.com/pdf/esx_SAN_guide.pdf

To configure the Emulex HBA with VMware ESX Server and Hitachi storage solutions, use the following procedure. Other Emulex parameters may be configured depending on the specific environment.

1. Reboot the server.
2. During system startup, press **ALT+E** to access the Emulex adapter settings menu.
3. Enter the Emulex HBA configuration utility during boot up and set these parameters:
   - Enter a number to select the HBA to be configured and press **Enter**.
   - From the **Adapter** menu, enter 2 to configure the adapter's parameters.
   - HBA Firmware: at the minimum level (3.90a7 for LP9002, 1.00a4 LP9802DC).
   - Host Adapter Bios: Disabled (option 1)
   - Topology Selection: Fabric Point-to-Point (option 4)
4. Repeat the configuration for other HBAs.
5. Press **X** to exit, save your changes, and reboot the server.

Chapter 3 Preparing for VMware ESX Server OS Configuration

This chapter describes how to prepare for VMware ESX Server OS configuration. Topics include:

- VMware installation (section 3.1)
- VMware configuration (section 3.2)
- Guest OS Preparation (section 3.3)

After preparing new storage device configuration, preparing for VMware ESX Server configuration is necessary before the Hitachi new storage device could be discovered and managed under VMware ESX Server environment.

Note: Configuring the VMware ESX Server for Hitachi TagmaStore AMS/WMS devices should only be performed by an authorized VMware system administrator and/or a Hitachi Data Systems representative. Configuration requires root access to the VMware ESX Server. If you have questions or concerns, contact the appropriate personnel and/or Hitachi Data Systems Support Center.

3.1 VMware Installation

Installing VMware ESX Server requires the following:

- VMware ESX Server Installation CD
- VMware ESX Server License keys
- VMware approved Intel server with fibre-channel HBA(s)
- A remote console with Internet Explorer 6.0 for accessing the VMware Management Interface to configure VMware ESX Server

A description of the VMware ESX Server installation can be found in VMware Installation Guide from VMware Web site. Booting from SAN requires a dedicate handling of vmhba usage as shared or non-shared device during the installation. However, VMware installation of booting from local device (internal) is simple and straight forward as these following steps:

1. Start VMware installation from a booting up VMware installation CD.
2. From the **Welcome Installer** window, type "text" for installation using text interface or press **Enter** to begin a standard installation using Graphical User-Friendly Interface method.

```
Boot: <enter>
```

3. Verify that VMware loading the internal SCSI driver and other appropriate driver (i.e., network driver, etc.) before loading the **GUI Installation** window

```
Loading aic7xxxx driver…
Loading e100 driver…
```

4. At the **GUI Welcome** window, click **Next** to proceed to the next page.
5. At **Installation Type**, select **Default Installation type** and click **Next**.
6. At the **End User License Agreement**, check the **I accept** box and click **Next**.
7. At the **VMware ESX Serial Numbers**, enter the appropriate licenses and click **Next**.
8. At **Device Allocation**, leave all options as default and click **Next**.
9. Verify that the following settings are part of the default setting:
   - There is at least one Ethernet controller is dedicated to Service Console.
   - At least one SCSI storage controller group (internal) is dedicated to Virtual Machine and shared with Service Console with **Shared with Service Console** box checked.
   - Fibre Channel Controller Group is dedicated to Virtual Machines with its HBAs are non-shared with Service Console.
   - The rest of the Ethernet controller should be dedicated to Virtual Machines.
8. At **Disk Partition Setup**, choose **Automatic Partitioning and Remove all Partitions** and click **Next**.

9. Verify that there is at least one internal drive discovered as sda device. For example:

```
sda: IBM DNES-318350Y SA60AKGo – 17359 MB
```

9. Click **Yes** to remove all partitions on /dev/sda.
10. At Disk Setup, verify that VMware has created the following partitions:
    - /dev/sda1 as /boot (boot device)
    - /dev/sda2 as / (root device)
    - /dev/sda3 as swap (swap space)
    - dev/sda4 as an extended partition

    We recommend that you delete the /dev/sda4 extended partition and leave it as free space, as it may be more flexible to manage it within the VMware Management Interface when creating several VMFS-2 volumes for different Guest OS.
11. At **Network Configuration**, enter the appropriate Hostname, network IP Address, and netmask, then click **Next**.
12. At **Time Zone Selection**, choose the appropriate time zone and click **Next**.
13. At **Account Configuration**, enter the password for the Root Password of the root user and click **Next**.
14. At **About to Install**, verify all corrected information has been provided and click **Next** to begin VMware ESX Server installation, or **Back** to go back for modification.
15. At **Installing Packages**, VMware installs several packages.
16. At **Congratulations**, verify the installation is complete, click **Next**, and remove the installation CD. The system will reboot.
17. At the boot prompt verify that the first entry in the **LILO** menu is the newly installed OS partition ESX Server, press Enter to boot ESX Server OS:

    Boot: <enter>
18. After the ESX Server boots, press **Alt-F2** to access to VMware Service Console.
19. Verify that VMware ESX Server 2.5.0 has loaded.

```
VMware ESX Server 2.5.0
Kernel 2.4.9-vmnix2 on an i386
VM4-DF700 login:
```

20. At the Login prompt and Password, enter the login ID and the root password.

The next step is to configure VMware ESX Server from the VMware Management Interface GUI.

3.2 VMware Configuration

The activities involved in configuring VMware for AMS/WMS are:

- Accessing VMware Management Interface (section 3.2.1)
- Verifying Startup Profile (section 3.2.2)
- Creating VMFS-2 volumes for Guest OS and Guest OS swap file (section 3.2.3)
- Configuring Swap File (section 3.2.4)
- Configuring networking (section 3.2.5)
- Configuring VMware to discover LUN-8 (section 3.2.6)

3.2.1 Accessing VMware Management Interface

From the Web browser of a remote console, point to the Hostname or the IP address of the ESX Server to access the VMware Management Interface. A security alert may be prompted before accessing the site. Proceed as normal or install the certificate where appropriate, as in the Local Computer of Trusted Root Certification Authorities by clicking **View Certificates**, and then proceed to **VMware Management Interface** login window.

Figure 3.1 Installing Certificate

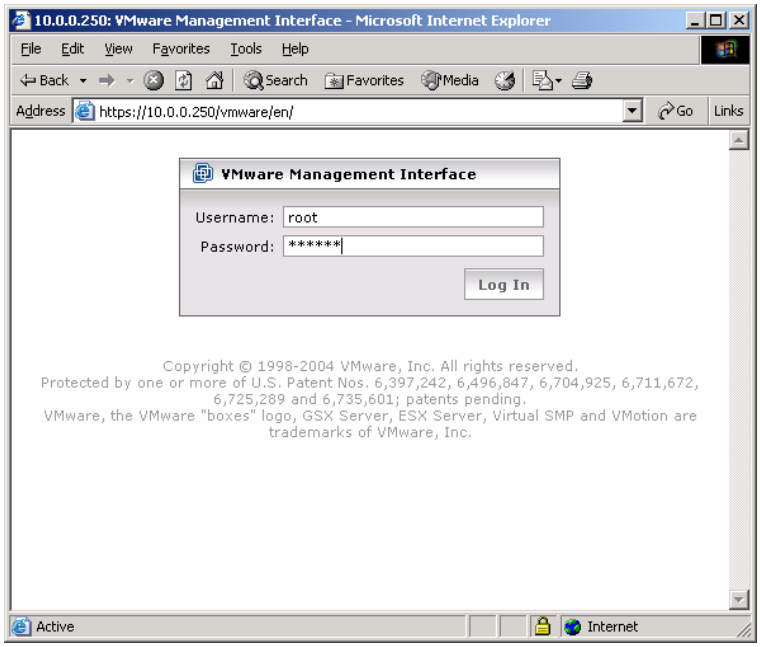

Figure 3.2 VMware Management Interface

3.2.2 Verifying Startup Profile

The startup profile contains VMware resources usage and assignment information about the startup profile that can be helpful when there are multiple virtual machine HBAs and internal SCSI adapters that use local boot and SAN boot configurations. Using the startup profile, you can recognize each virtual device.

Figure 3.3 shows a typical startup profile. In this example, there are several virtual Ethernet controllers, one dedicated to the service console, one internal SCSI virtual controller, and several virtual external fibre-channel HBAs. Memory reserved for service console is set at default level (maximum 8 Guest OS virtual machines).

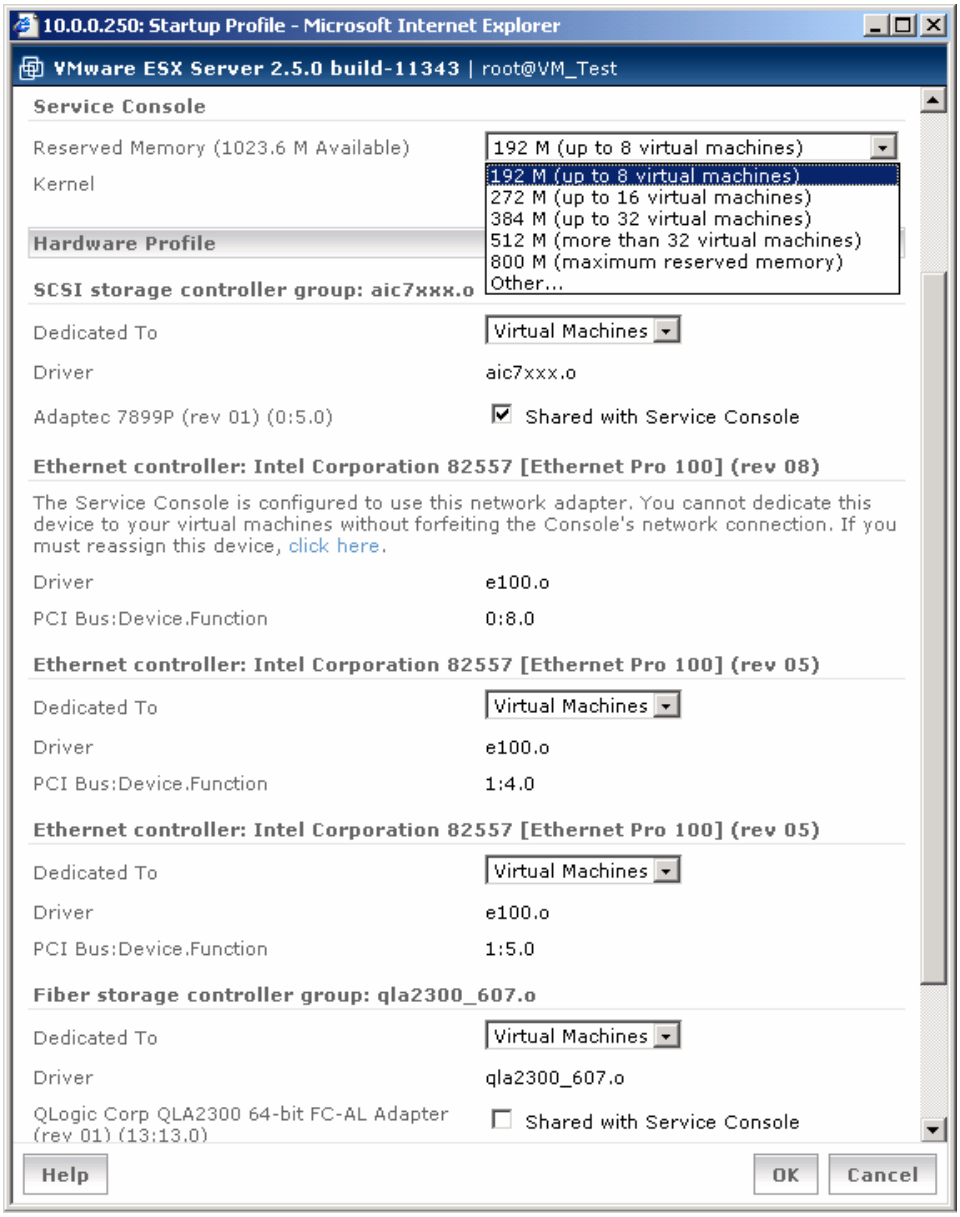

Figure 3.3 VMware Startup Profile

3.2.3 Creating VMFS-2 Volumes for Guest OS and Guest OS Swap File

Typically, the server has one internal drive that holds everything (OS, data, swap file, etc.). For improved manageability, more than one internal drive can be used to separate the OS and the data. In this scenario, one internal drive is dedicated for VMware OS, while the other is dedicated for the Guests OS virtual machine and swap file. In this example, there are several Guest OS virtual machines, VMFS-2 partition, and swap file in one extended partition on the single internal drive.

To create a VMFS-2 volume for Guest OS and Guest OS swap file, access Disks and LUNs of Storage Management under the options for VMware Management Interface and recognize the internal virtual machine HBA vmhba0:x:y with three native Linux partitions and a free disk space.

1. Click **Create Volume** of the free disk to create a partition (see Figure 3.4).
2. Choose **Custom** as the creating method (see Figure 3.5) and click **Next**.
3. For the file system type, choose **Extended Partition** (see Figure 3.6) and click **Finish**.
4. Verify from the **Disks and LUNs** menu that an extended partition has been created with the entire partition size as free space (see Figure 3.7).
5. Since there is an extended partition, you can create more than four VMFS-2 volumes for the Guest OS and swap device. Click **Create Volume** on free space to create VMFS-2 volumes.
6. Click **Custom** again and select VMFS-2 file system for a new VMFS-2 volume.
7. Enter Volume Label with Local Guest OS – Windows 2003 and Capacity 3500 MB. Leave **Maximum File Size** and **Access Mode** at their default settings of 456 GB and public access, respectively (see Figure 3.8).
8. Click **Finish**.

   Note: Consult your Guest OS vendor for appropriate disk partition size for OS installation.

9. From the **Disks and LUNs** menu, verify that Windows 2003 Guest OS VMFS-2 has been created. For example:

```
5. Local Guest OS – Windows 2003 (Logical VMFS-2.11 Volume)
        Capacity              3.42 G
        Maximum File Size     456.00 G
        Access Mode           Public
```

10. Repeat creating VMFS-2 volumes for other Guest OS virtual machines.
11. Create a VMFS-2 volume for Guest OS swap file with the label Guest OS Swap and a recommended volume capacity of 2GB (see Figure 3.9).

    Note: The capacity you specify depends on how much the total system memory is and each Guest OS virtual machine memory allocated.

For detailed instructions about how to create a VMFS-2 volume, refer to section 4.3.

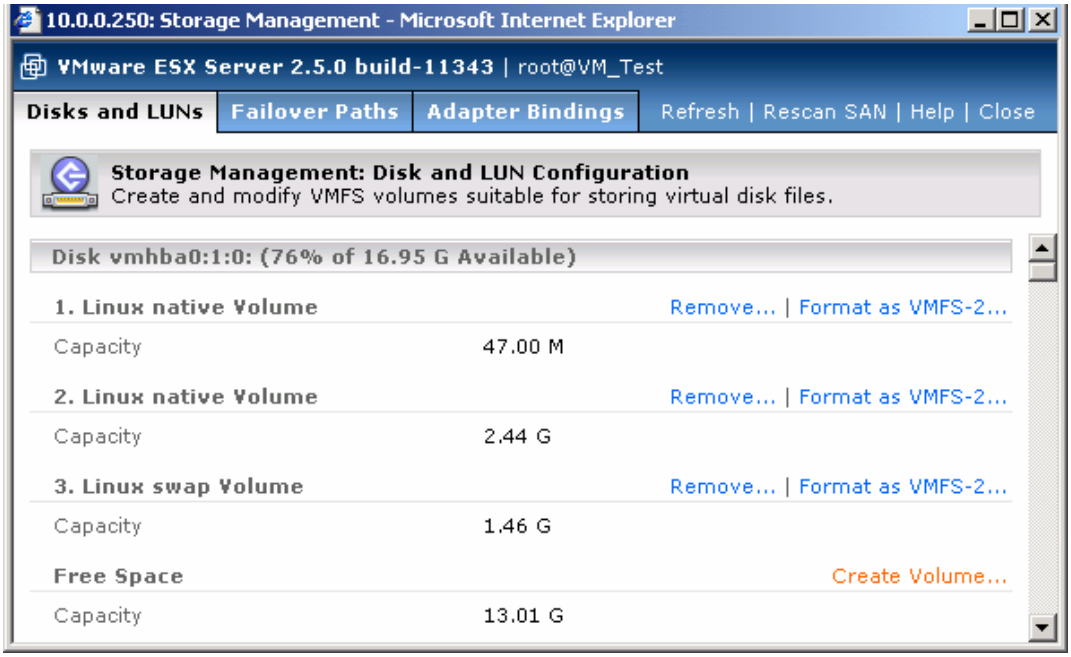

Figure 3.4 Creating Volumes for Free Space

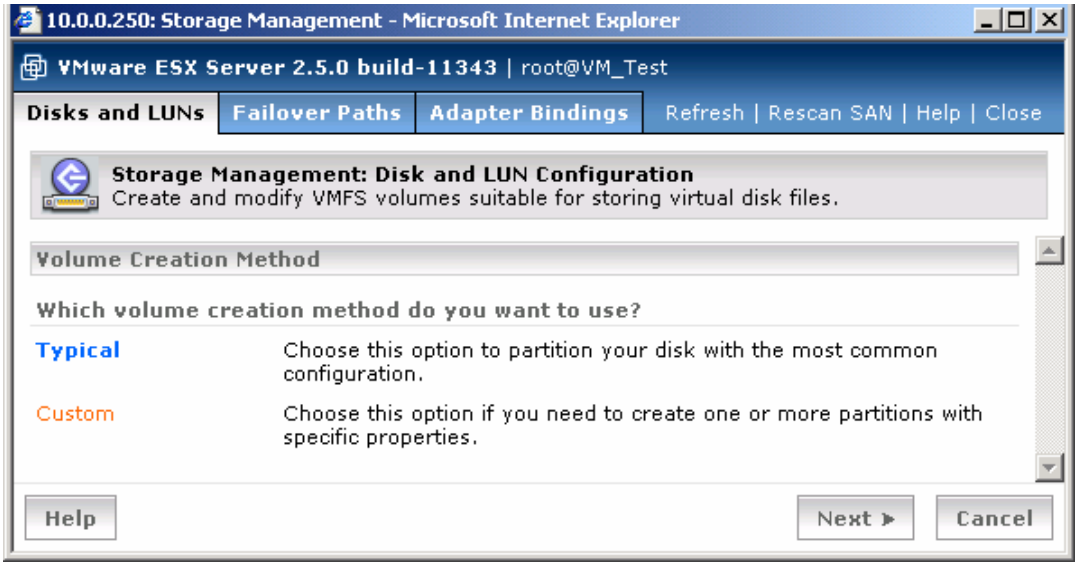

Figure 3.5 Choosing Custom Method

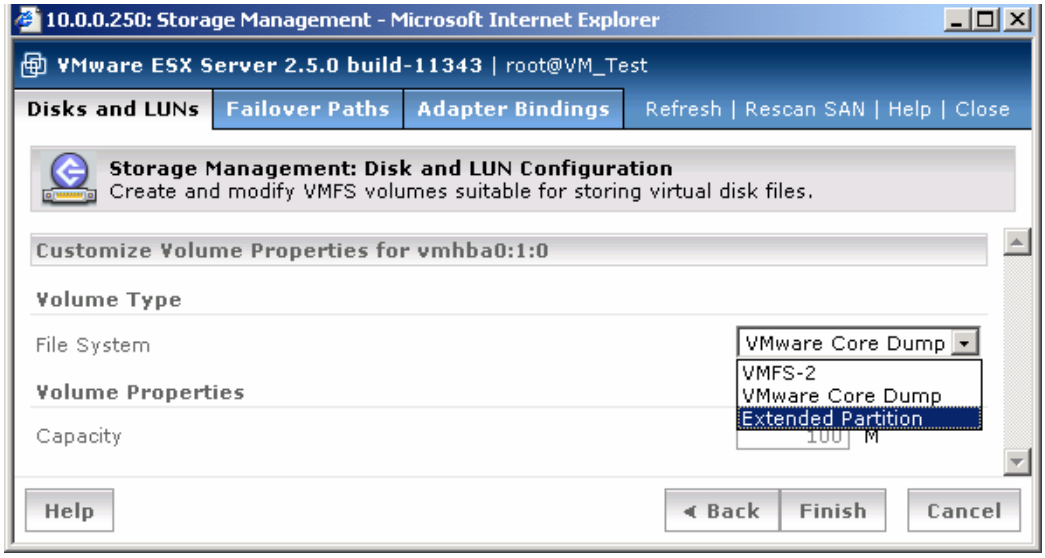

Figure 3.6 Selecting Extended Partition

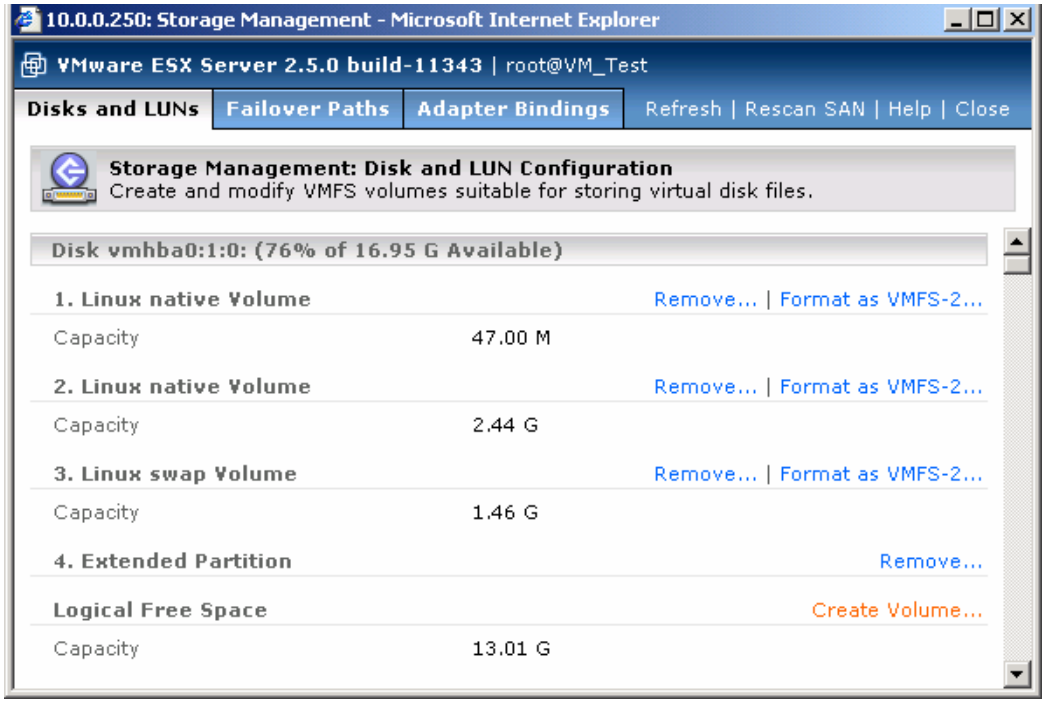

Figure 3.7 Extended Partition Created

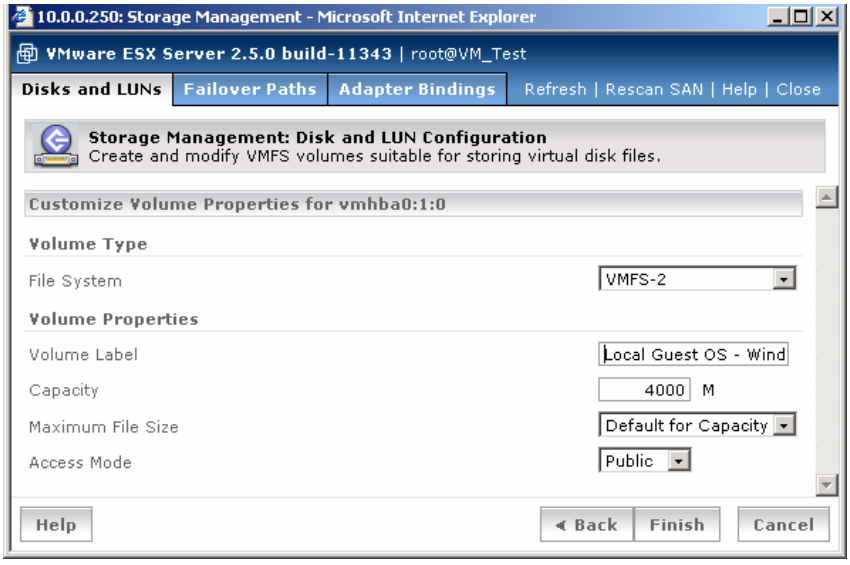

Figure 3.8 Creating VMFS-2 Volume for Guest OS

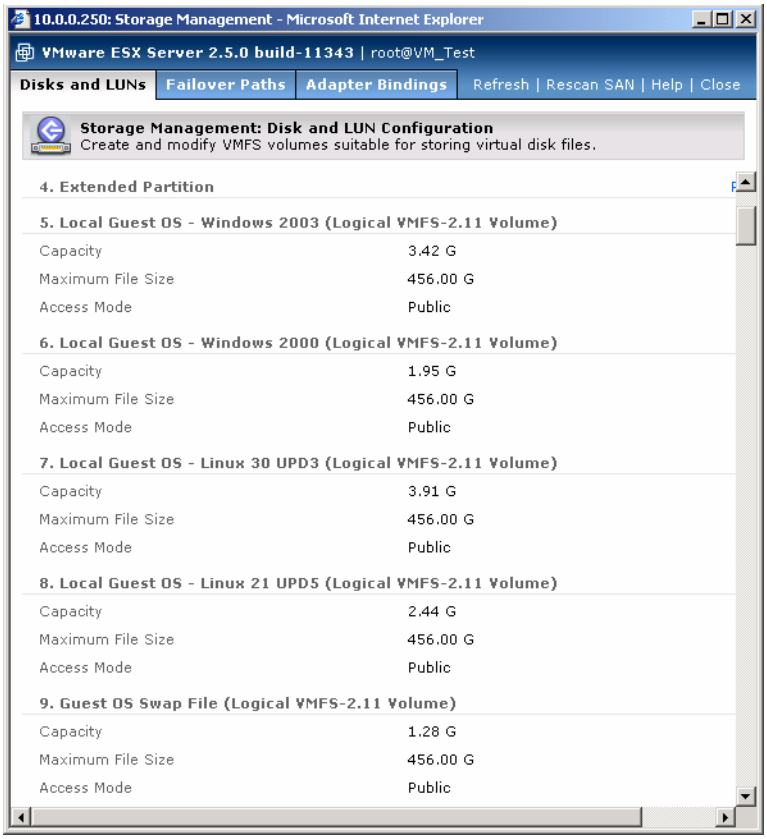

Figure 3.9 Guest OS and Swap Partitions

3.2.4 Configuring a Swap File

Depending on the amount of memory that the VMware server can allocate to the virtual Guest OS, a Guest OS swap file may be necessary. A swap file is particularly helpful during periods of high usage and at startup times for multiple Guest OS virtual machines.

Initially, the swap device can be configured on the Status Monitor page of the VMware Management Interface (see Figure 3.10).

1. Click **Reconfigure** to configure the Guest OS swap file.
2. At the **Swap Configuration** window, select the VMFS-2 volume labeled Guest OS Swap File and enter the maximum available file size (see Figure 3.11). The default file name is Swapfile.vswp and activation policy is active at system startup.
3. Click **OK** to create the Guest OS swap file (see Figure 3.12).
4. Reboot the server for the configuration to take effect.

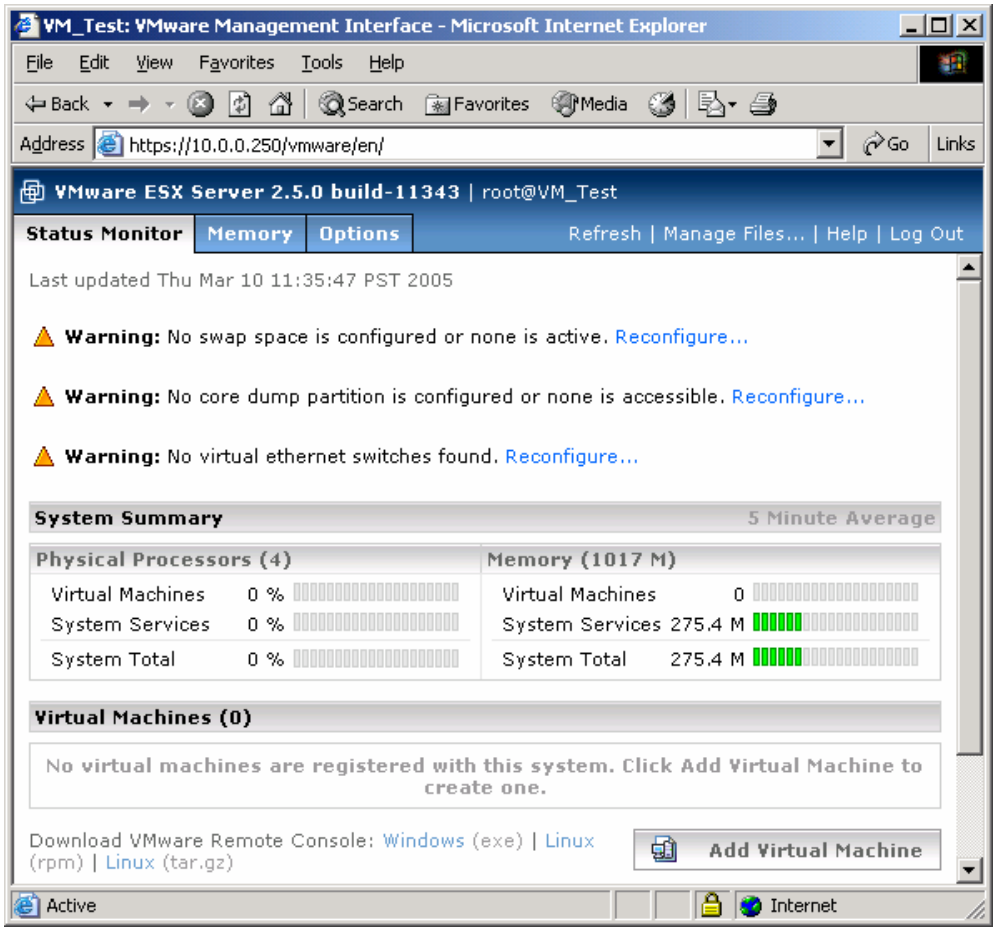

Figure 3.10 Reconfiguring the Swap File

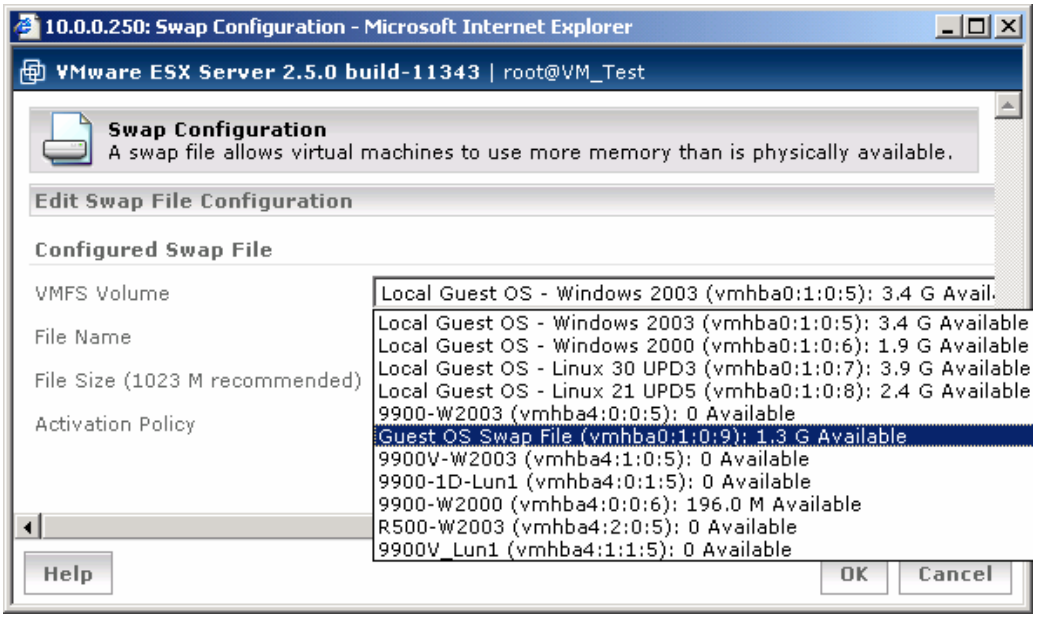

Figure 3.11 Selecting VMFS-2 Guest OS Swap File

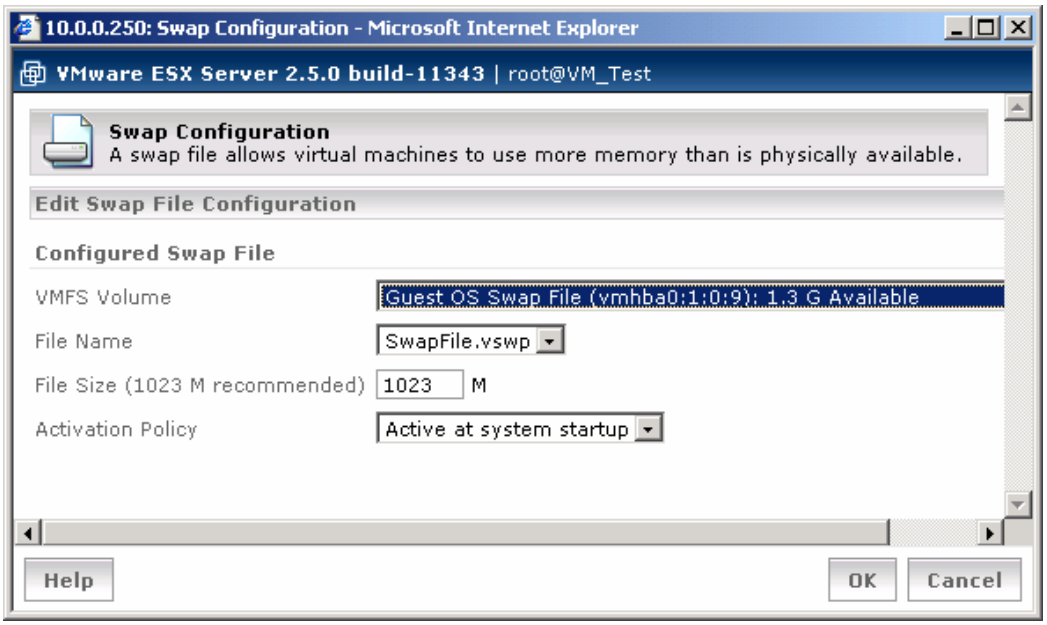

Figure 3.12 Guest OS Swap File Created

3.2.5 Configuring the Networking

There is no need to configure networking as long as all Guest OS virtual machines operate independently as VMware virtual devices. However, for the Guest OS virtual machines to communicate with each other and with the VMware ESX Server, a virtual networking switch must be created. A virtual networking switch can also be bound to a physical networking adapter so that the Guest OS virtual machines could communicate outside of its virtual network (i.e., another physical server).

1. When a warning message indicates that you have not configured your networking, click **Reconfigure** to configure the networking as in Reconfiguring Swap file.

   Note: Networking configuration can also be done under **Network Connections of VMware Management Interface Options**.

2. Click **Create** to create a virtual switch. The page in Figure 3.13 appears.
3. Enter the network label for the virtual networking switch, such as **Virtual – Network 0**.
4. Apply on the outbound adapter box of a connected Ethernet adapter to configure outbound communication.
5. Click **Create Switch** to create the virtual switch (see Figure 3.14).

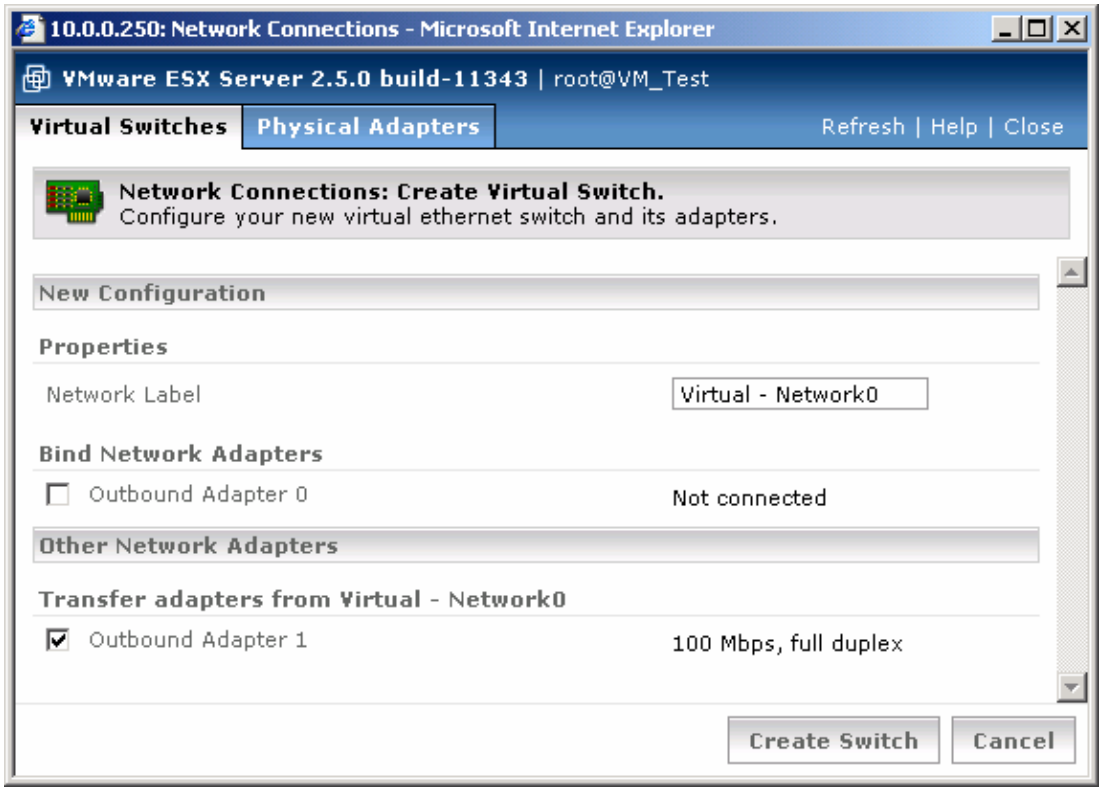

Figure 3.13 Configuring Virtual Switch

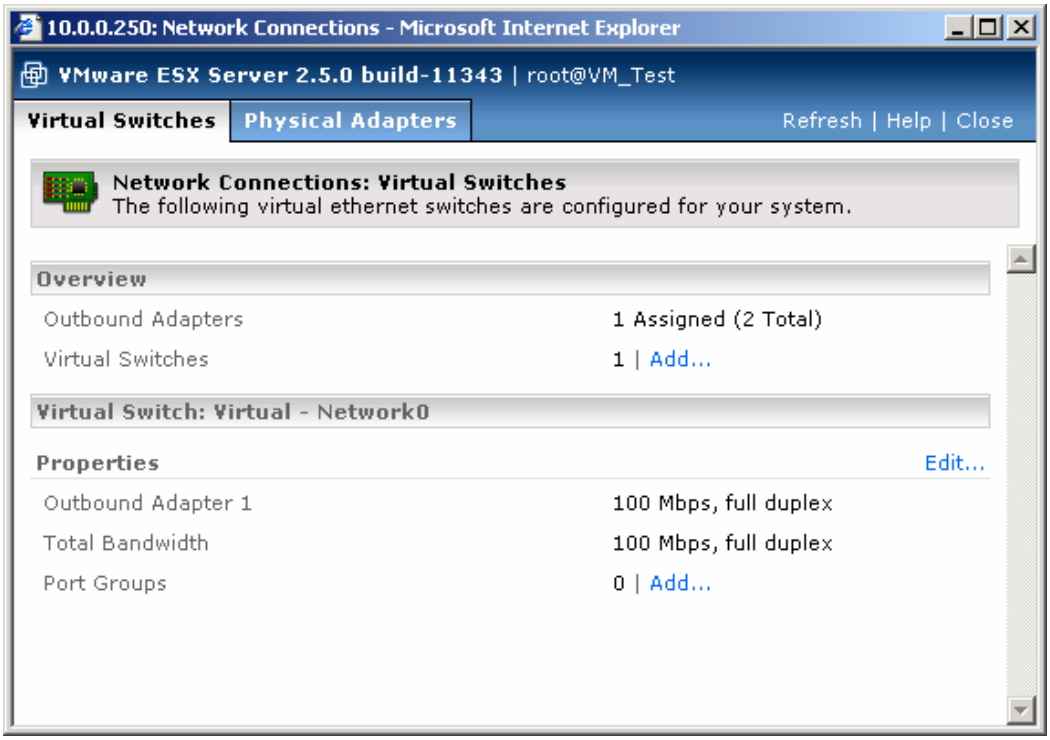

Figure 3.14 Virtual Switch Created

3.2.6 Configuring VMware to Discover Lun8

There are several VMware parameters you can set to fine-tune the operation of the VMware ESX Server and its Guest OS virtual machines. These settings reside under the Advanced Settings of the VMware Management Interface Options (see Figure 3.15).

One parameter, Disk.MaxLun, determines the maximum number of Luns that VMware can discover. The default Disk.MaxLun value (8) lets the ESX server discover Lun0 to Lun7. If your environment uses Lun8, change the Disk.MaxLun setting to 9. You must make this change even if your setup uses a single Lun configured as Lun8. If you do not change the default value for Disk.MaxLun, VMware will not discover Lun8.

The following example shows how to specify this parameter under Advanced Settings of the VMware Management Interface Options (also, refer to Figure 3.16 and Figure 3.17):

Disk.MaxLUN = 40 (or any appropriate number)

Note: You must be an authorized VMware system administrator to alter the parameter. VMware supports maximum 128 devices per array, per system, per server

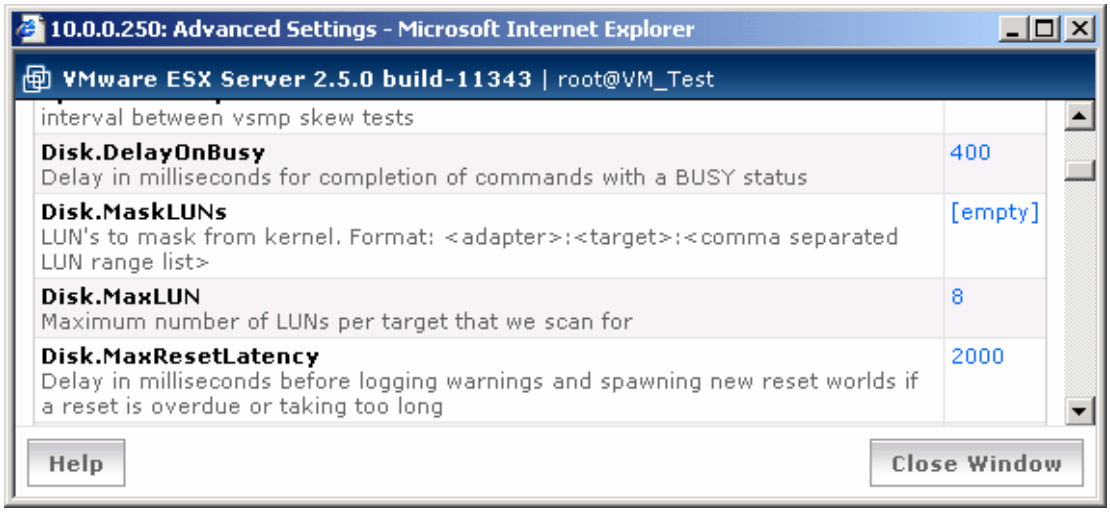

Figure 3.15 Disk.MaxLUN of Advanced Settings

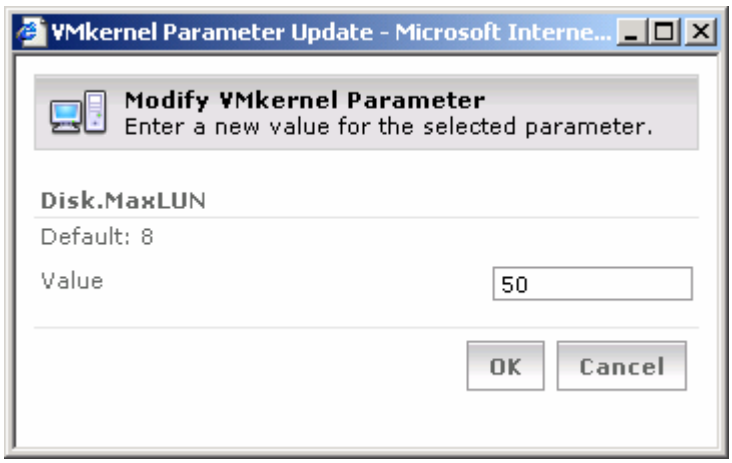

Figure 3.16 Modifying Disk.MaxLUN Value

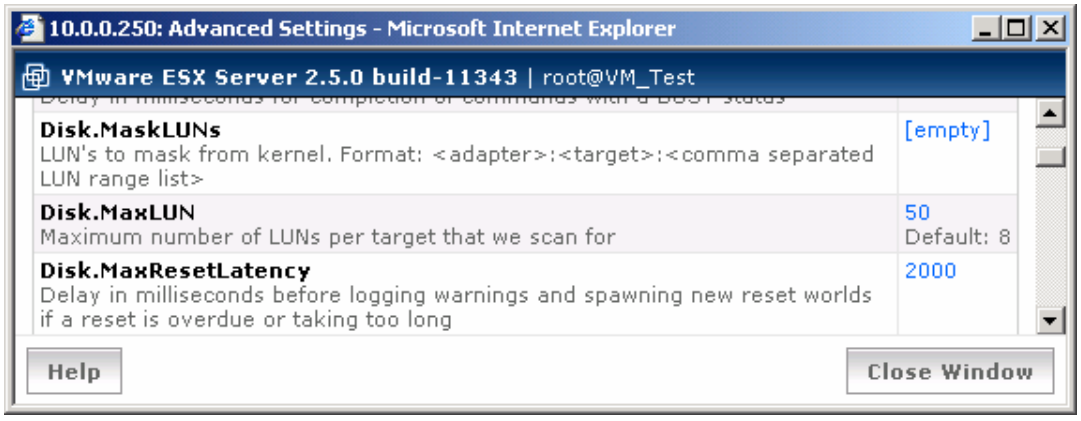

Figure 3.17 Disk.MaxLUN Value Modified

3.3 Guest OS Preparation

VMware treats all physical components of the ESX Server as virtual devices. The Guest OS, in this sense, is also a virtual device. Each Guest OS resides in a VMKFS-2 file managed by VMware OS. Like any other VMKFS-2 file, Guest OS VMKFS-2 file can be moved, added, deleted, or copied to different places.

The activities involved in preparing VMware Guest OS for Hitachi TagmaStore AMS/WMS are:

- Creating and identifying VMKFS-2 volume for VMware Guest OS (see section 3.3.1)
- Creating VMX folder (see section 3.2.2)
- Installing Guest OS (see section 3.2.3)

3.3.1 Creating and Identifying VMKFS-2 Volume for VMware Guest OS

Once the VMFS-2 volume for the VMware Guest OS has been created (see section 3.2.3), you can verify that the required Guest OS VMFS-2 has been created with the appropriate labeling by:

- Using the Disks and LUNs section of the VMware Management Interface (see Figure 3.18).
- Accessing VMware File Management. Click **Manage Files** from VMware Management Interface and access the vmfs directory for the entry (see Figure 3.19).

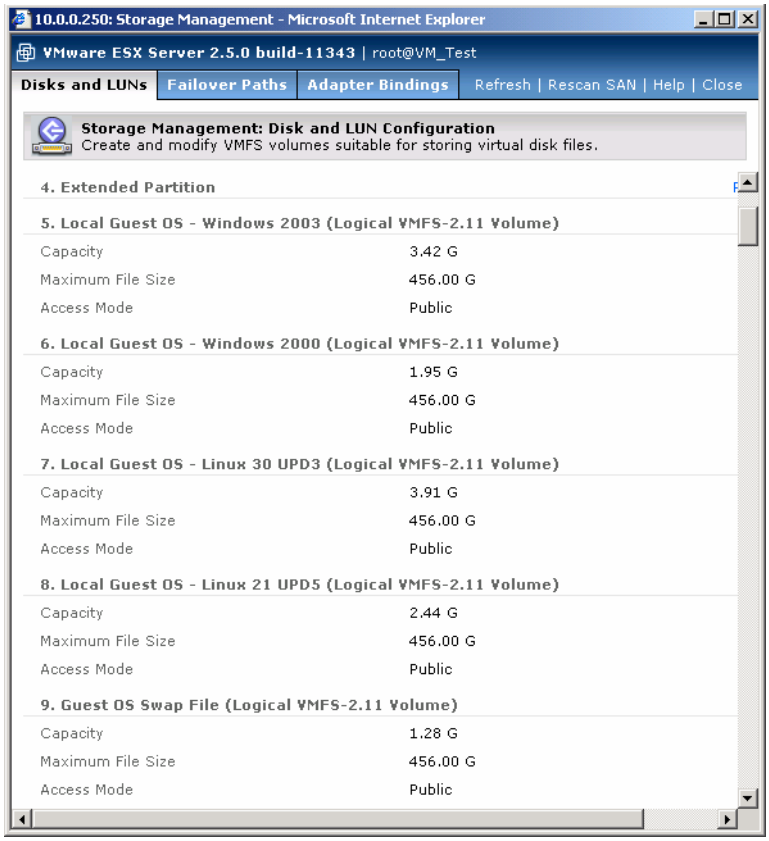

Figure 3.18 Identifying Guest OS VMFS-2 Volumes

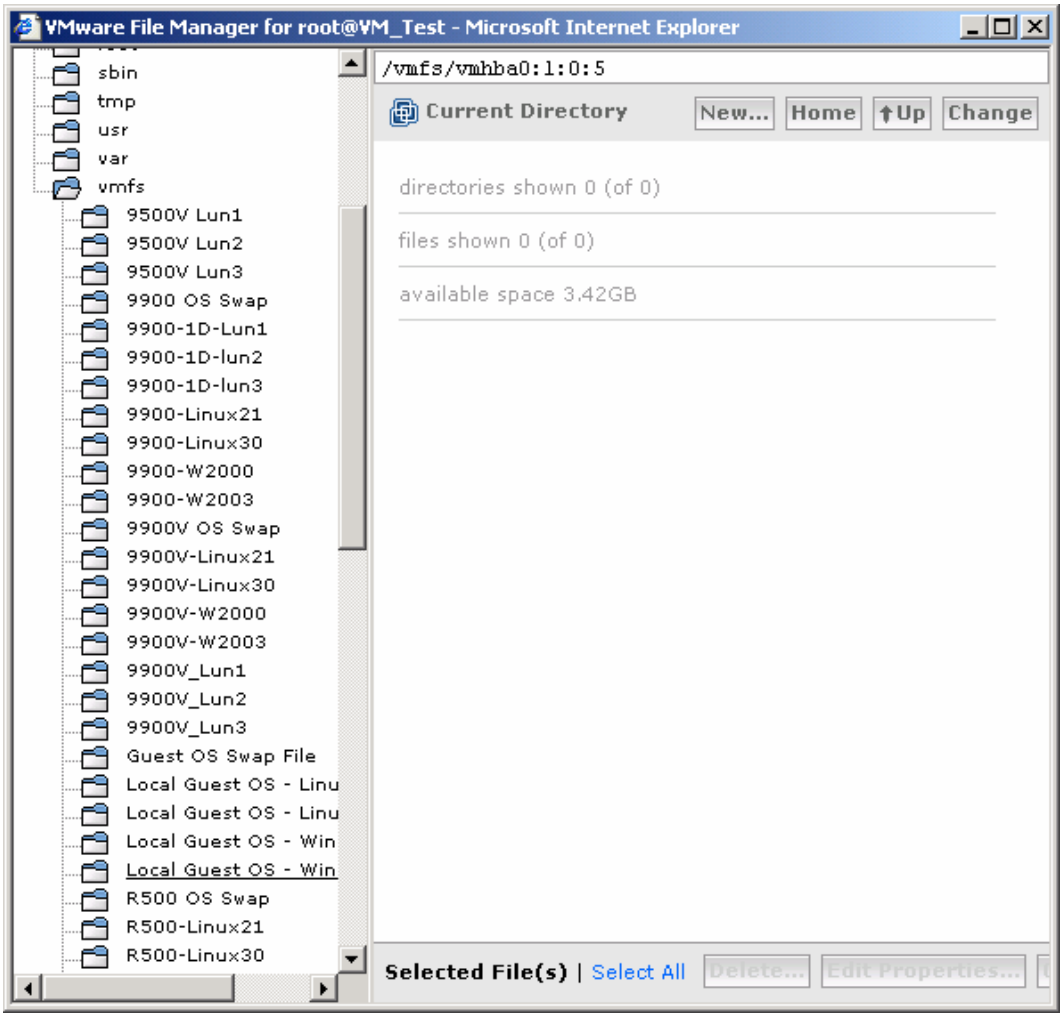

Figure 3.19 VMware File Manager for Guest OS VMFS-2 Volume

3.3.2 Creating a VMX Folder

Once the VMFS-2 volume has been created and identified, a vmx file is created for the Guest OS. This file will hold all relevant information about the Guest OS. By contrast, the VMFS-2 volume contains the operating systems of the Guest OS (i.e., Windows 2003, Linux 2.1).

1. From the **Status Monitor of VMware™ Management Interface**, click **Add Virtual Machine** (see Figure 3.20).
2. Select a Guest OS to be created and label the Guest OS display as desired (see Figure 3.21).

   Note: The default vmx name is winNetEnterprise.vmx for Windows 2003 EE, win2000AdServ.vmx for Windows 2000 AS, and linux.vmx for all Linux versions.
3. Change the location of the vmx file as appropriate (i.e., winNetEntetrprise-Sys1.vmx) or accept the default.
4. Click **Next**.

5. Choose the number of virtual processors (1 or 2) and virtual memory allocation for the Guest OS (suggested or minimum) or size according to the total system memory availability (see Figure 3.22).
6. Click **Next**.
7. From the Virtual Disk type selection, choose **Blank** as a newly created Guest OS (see Figure 3.23).
8. Select the image file location for the Guest OS (i.e., Local Guest OS – Windows 2003,). Enter a name for the image vmdk file name (i.e., LocalW2003EE,) change or leave the capacity as default, and leave virtual **SCSI mode** and **Disk Mode** at default value, 0:0 and Persistent options (see Figure 3.24).
9. Click **Next**.
10. At the **Hardware Configuration** window, verify that Processors, SCSI Controller 0, and Virtual Disk (SCSI 0:0) VMware disk Image created successfully. Other removable devices such as floppy, CD drive, and network adapter should be available and connected at power on. Click **Close** to close the **Hardware** menu.
11. At the **Status Monitor** menu, a Guest OS virtual machine appears with the appropriate label (see Figure 3.25).

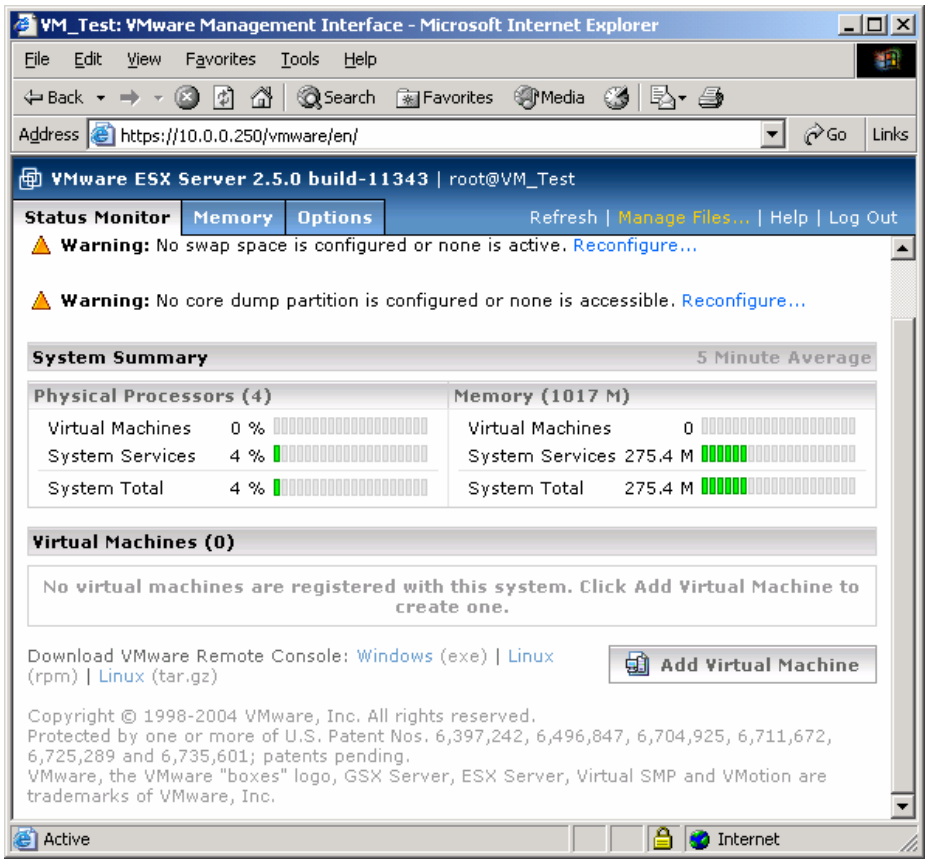

Figure 3.20 Add Virtual Machine

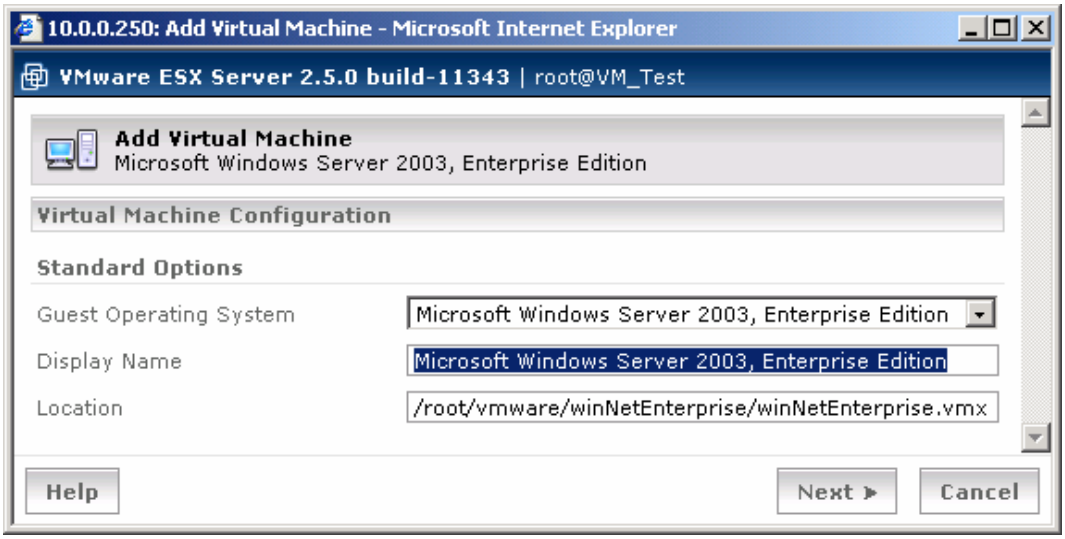

Figure 3.21 Selecting and Naming a VMX file

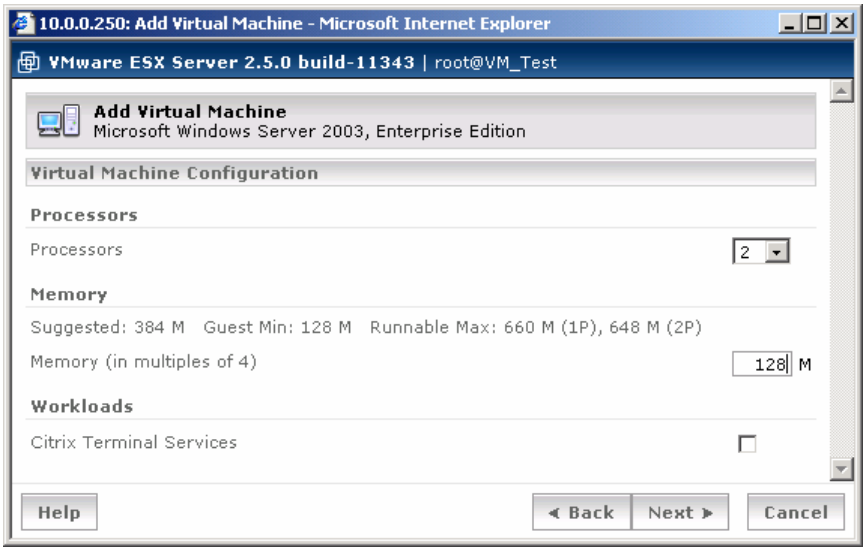

Figure 3.22 Virtual Processors and Memory Allocation for Guest OS

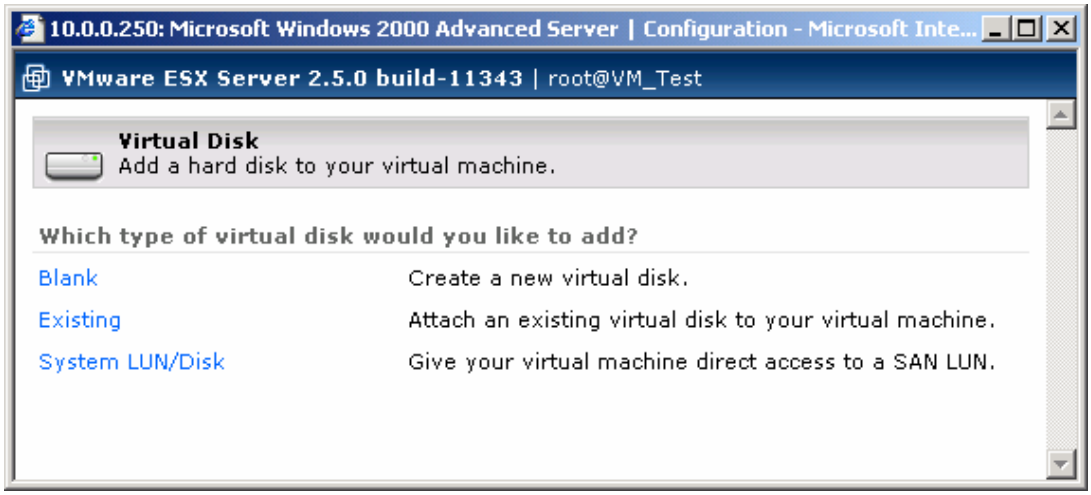

Figure 3.23 Virtual Disk Type Selection for Guest OS Installation

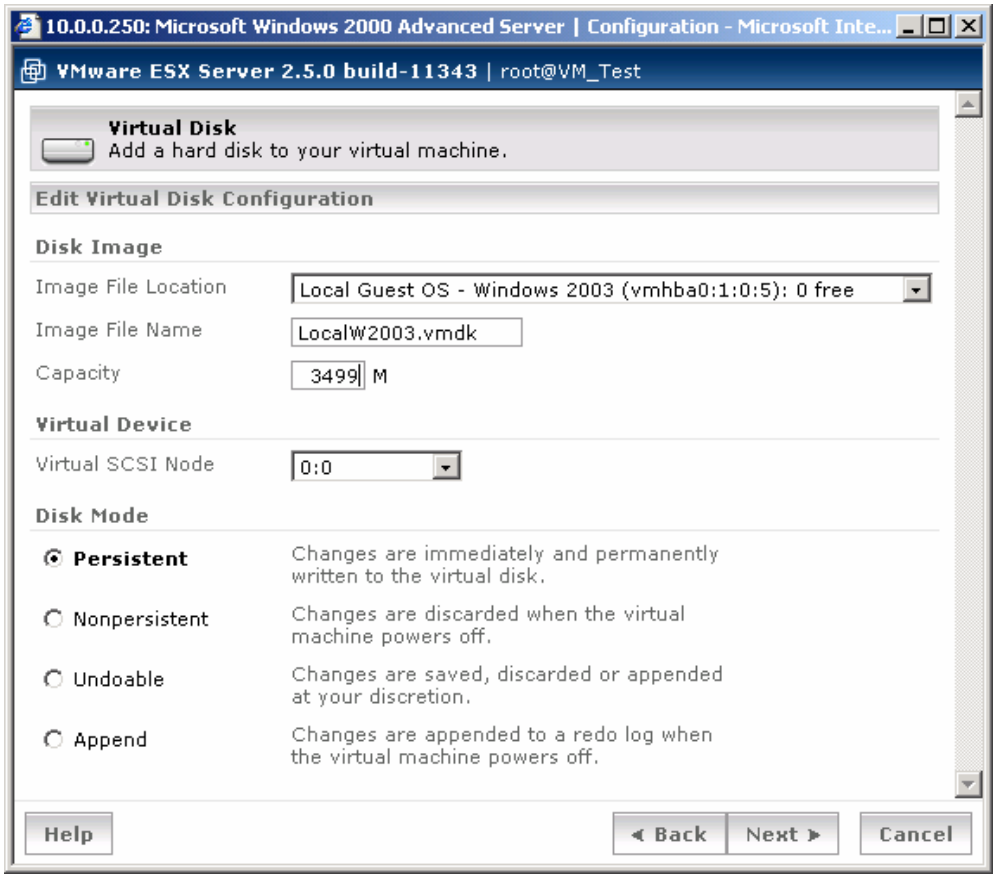

Figure 3.24 Configuring the Guest OS Image File

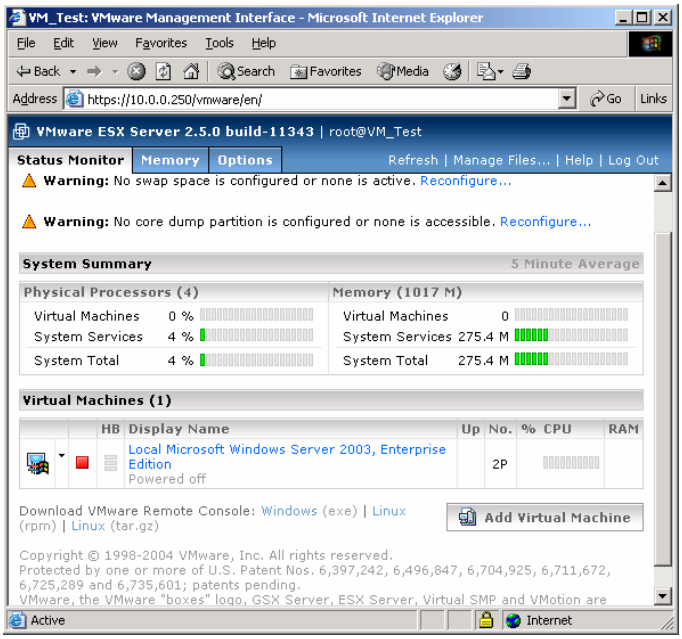

Figure 3.25 Guest OS VMX File Created

Figure 3.26 shows the windows for verifying the Guest OS image file vmdk through VMware File Manager.

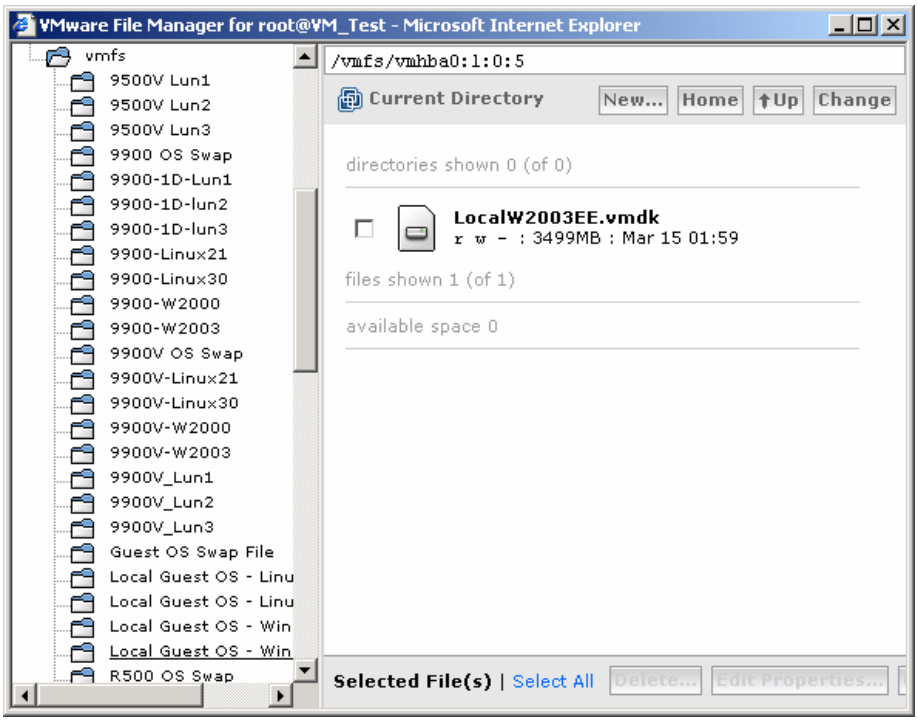

Figure 3.26 VMware File Manager for Guest OS vmdk Image File Verification

3.3.3 Installing the Guest OS

To install an operating system onto a Guest OS virtual machine:

1. Download and install VMware remote console from the Status Monitor of VMware Management Interface (see Figure 3.27).
2. Place the installing Operating System CD in the ESX Server CD-ROM drive.
3. Power on the Guest OS machine by clicking on the red square button and then the green arrow icon (see Figure 3.28).
4. Double-click the terminal icon to attach a remote console to the Guest OS virtual (see Figure 3.29).
5. Open the remote console and enter the appropriate root password to access the VMware ESX Server, then click **Connect** (see Figure 3.30).
6. Verify VMware is installing the Operating System for the Guest OS through the remote console (see Figure 3.31).

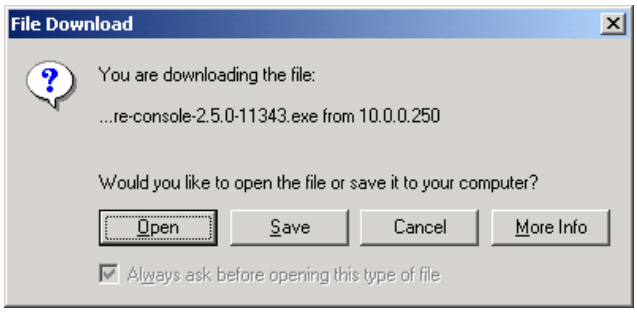

Figure 3.27 Downloading VMware Remote Console

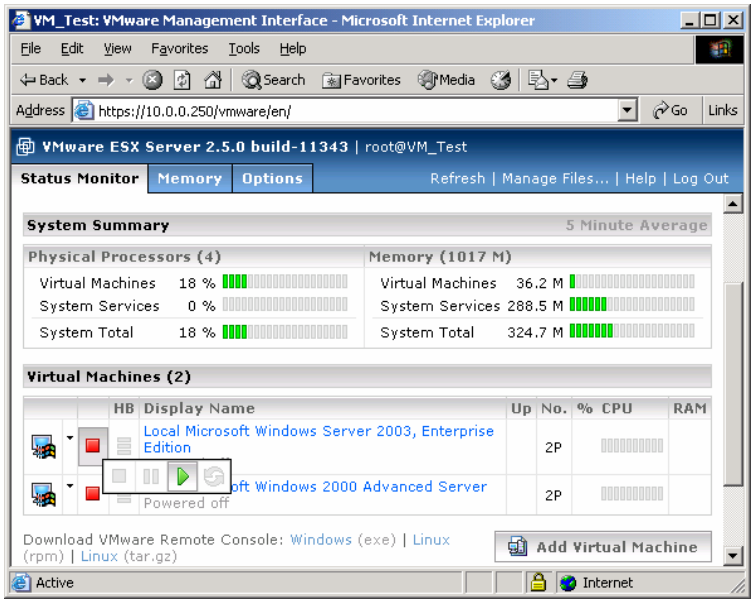

Figure 3.28 Power-on VMware Guest OS Virtual Machine

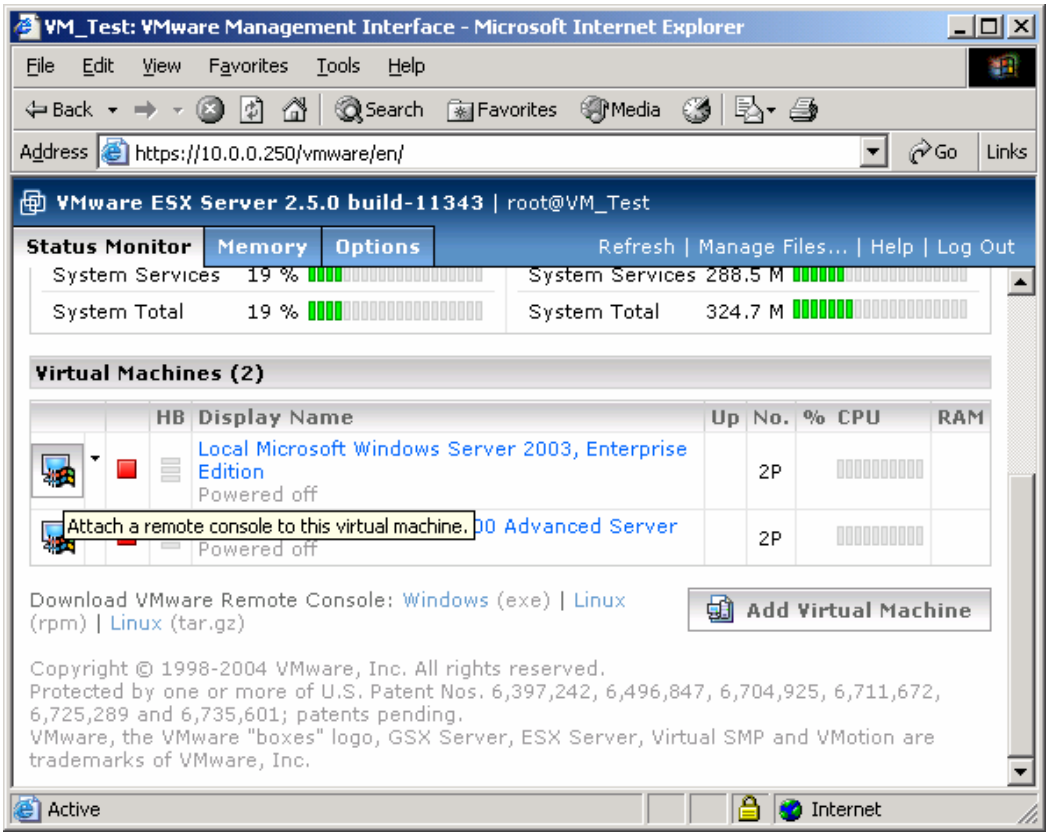

Figure 3.29 Attaching a Remote Console to Guest OS Virtual Machine

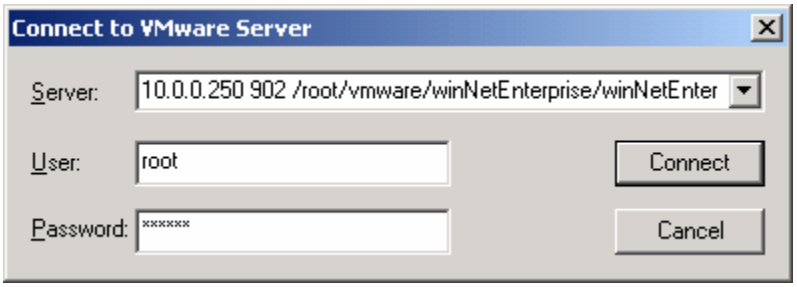

Figure 3.30 Connecting to the VMware ESX Server

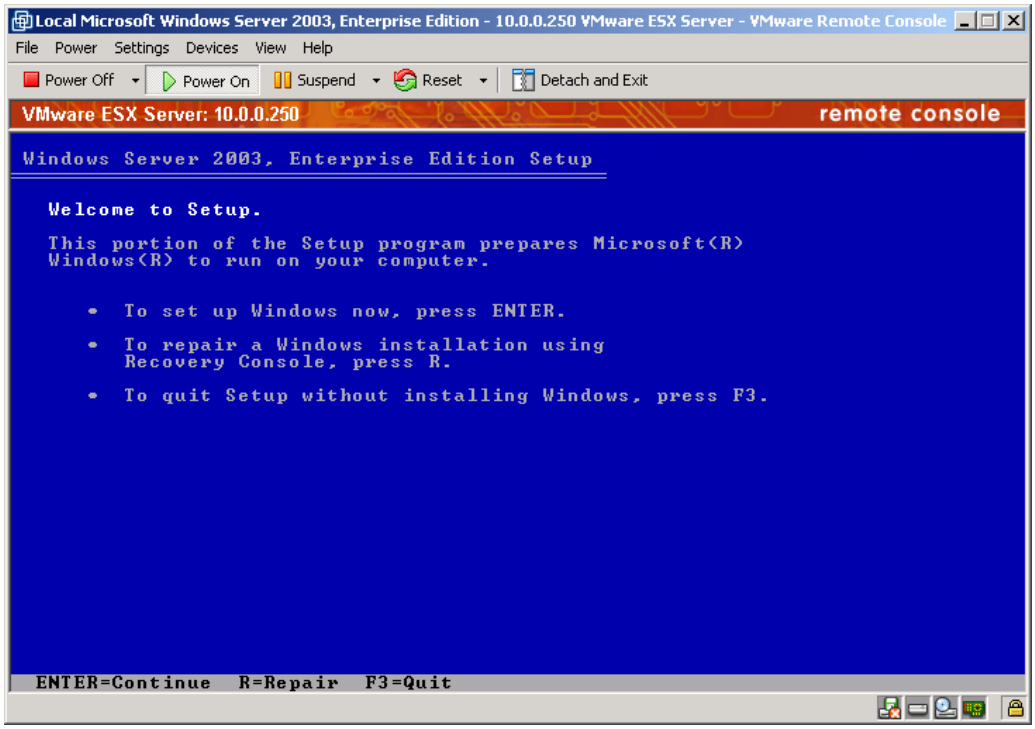

Figure 3.31 Guest OS Virtual Machine Operating System Installation

Chapter 4 Configuring the New Devices

This chapter describes how to configure the TagmaStore AMS/WMS devices. This process involves:

- Verifying the VMware New Storage Device (section 4.1)
- Creating the VMware VMFS-2 Volume (section 4.2)
- Assigning VMFS-2 and Raw Devices to the Guest OS (section 4.3)
- Managing the Guest OS (section 4.4)
- Verifying Guest OS access (section 4.5)
- Online LUN addition and deletion (section 4.6)

After the TagmaStore AMS/WMS installation procedure has been completed and new device recognition has been verified, the devices on the newly installed TagmaStore AMS/WMS system can be configured for use.

Note: Configuration of the TagmaStore AMS/WMS devices should only be performed by a VMware system administrator and/or a Hitachi Data Systems representative. Configuration requires root access to the host system. If you have questions or concerns, contact the Hitachi Data Systems Support Center.

Note on the term "SCSI disk": The TagmaStore AMS/WMS logical devices are defined to the host as SCSI disk devices, even though the interface is fibre-channel.

4.1 Verifying VMware New Storage Device

The first step in configuring the newly installed storage devices is to verify that the VMware system recognizes the new devices. Display the device information using the `ls` and `cat` commands on the ESX Server service console and verify that the ESX Server recognizes the devices. Then record each vmhba number and SCSI target ID (TID) for storage port to be used during later VMware volumes creation, labeling, and Guest OS LUNs assignment. Figure 4.1 shows a sample output of the `ls` command used to verify discovered LUNs of the storage attached to a vmhba (i.e., vmhba2.) Table 4.1 shows a sample worksheet for recording the vmhba number, storage target SCSI ID assignment, and LUN numbers.

1. Restart the ESX Server connected to the storage port. After rebooting, the ESX Server™ scans each HBA port attached to each storage port and assigns the appropriate vmhba number and SCSI TID. Entry for each discovered LUN is placed under the directory /proc/VMware/scsci/vmhbax, where X is the vmhba number assigned for each HBA FC port.
2. After the ESX Server reboots, login ESX Server and verify HBA log in and LUN discovering.
3. Verify that the log-in (green) LED is ON at the back of the two HBAs.
4. Verify that each HBA and internal SCSI controller is discovered as a VMware HBA with a SCSI ID number. Issue the command `cd /proc/VMware/scsi` and verify , for example:

   ```
   vmhba0    vmhba1    vmhba2    vmhba3    vmhba4    vmhba5 ……
   ```

5. Assuming that vmhba0 and vmhba1 correspond to the internal SCSI controllers, verify that the primary fibre channel HBA, vmhba2, discovers all LUNs assigned to it. Issue `ls /proc/VMware/scsi/vmhba2` and verify the following reply:

   x:0 x:1 x:2 x:3 x:4 ……..

   where x is the target SCSI ID number for a target port attached to a host bust adapter (vmhba2.) 0,1,2,3,4 are the LUN number assigned to that target port. Set the parameter **Disk.MaxLUN** from the **Advanced Setting** on the **Options** tab of the VMware Interface Management to an appropriate number if the discovered LUN is greater than Lun7. Rebooting ESX Server may be necessary to have the change in effect.
6. Assuming that vmhba2 is the primary and vmhba3 the secondary path of VMware multi-path function, issue `ls /proc/VMware/scsi/vmhba3` to discover any LUN entries. The display should be:

   stats
7. Verify that each LUN has two paths. Issue `cat/proc/VMware/scsi/vmhba2/0:0` on the primary path vmhba2 for Lun0 of SCSI TID 0. Verify LUN paths in the form of:

   vmhba2:x:0 on*#

   vmhba3:y:0 on

   where X is the target SCSI ID for the primary path vmhba2 and Y is the secondary path vmhba3. Repeat the command and verify both paths for every discovered LUN.

   Note: The `cat` command on the LUN of secondary path vmhba3 should only return "No such file or directory."

Complete LUN discovery and path failover information can also be verified with the VMware Management Interface under submenu **Option/Storage Management/Disks and LUNs**, and **Option/Storage/Management/Failover Paths**.

```
[root@VM4 root] ls /proc/VMware/scsi/vmhba2
0:0  0:2  0:4  0:6  0:8
0:1  0:3  0:5  0:7  0:9
0:10
[root@VM4 root] ls /proc/VMware/scsi/vmhba4
1:0  1:2  1:4  1:6  1:8
1:1  1:3  1:5  1:7  1:9
1:10
[root@VM4 root] cat /proc/VMware/scsi/vmhba2/0:0
Paths:fixed
     vmhba2:0:0 on*#
     vmhba3:0:0 on
```

Figure 4.1 Verifying New Device Recognition

Table 4.1 Sample SCSI Target ID Information Worksheet

Physical HBA	VMHBA# (Path)	SCSI Target ID (FC Port)	LDEV:LUN#	VMware™ SCSI ID Assignment
1st QLA2342 port 1	2 (Primary)	0 (port 1A)	0:0	Vmhba2:0:0
			0:1	Vmhba2:0:1
			0:2	Vmhba2:0:2
			…	…
1st QLA2342 port 2	3 (Secondary)	0 (port 2A)	0:0	Vmhba3:0:0
			0:1	Vmhba3:0:1
			0:2	Vmhba3:0:2
			0:3	Vmhba3:0:3
			…	…
2nd QLA2342 port 1	4 (Primary)	1 (port 1B)	2:0	Vmhba4:0:0
			2:1	Vmhba4:0:1
			2:2	Vmhba4:0:2
			2:3	Vmhba4:0:3
2nd QLA2342 port 2	5 (Secondary)	1 (port 2B)	2:0	Vmhba5:0:0
			2:1	Vmhba5:0:1
			2:2	Vmhba5:0:2
			2:3	Vmhba5:0:3
and so on….				

Note: In Table 4.1, ports 1A and 2A share the same LUNs, and ports 1B and 2B share the same LUNs. Also, the SCSI TID for ports 1B/2B can be changed to 0 (instead of 1).

4.2 Creating the VMware VMFS-2 Volume

After verifying new device recognition, create VMFS-2 Volumes on the new disk devices. This section describes the volume VMFS-2 creation process. For VMware ESX Server 2.5, refer to Appendix A for instructions on using the VMware Interface Management tool to perform various VMware configuration and device management.

Note: Do not create a VMFS-2 on a disk device that will be accessed as a raw device (e.g., some database applications use raw devices.)

To perform the virtual machine LUN assignment procedure:

- Log in to the VMware Management Interface.
- Verify all discovered LUNs.
- Prepare and format the LUNs with appropriate file systems.
- Assign LUNs to VMware virtual machines (section 3.3).

Perform these steps under the Storage Management option of VMware Management Interface. In addition, verify the Failover Path information under **Storage Management** for available failover paths.

1. From the **Options** menu of VMware Management Interface, open the **Disks and LUNs** tab under **Storage Management** (see Figure 4.2). Verify that the active primary vmhba is present and that each one has the correct number of attached LUNs. Each LUN is displayed in the following format:

   Disk vmhbax:y:z

   Note: Identify the internal vmhba and its attached LUN (local), as well as the external fibre channel vmhba with its discovered LUNs. In most cases, the internal vmhba appears as vmhba0 and vmhba1.

2. After identifying all required LUNs, click **Remove…** on the far right (see Figure 4.3). This will let you delete the existing partition on the LUN whose file system is different than VMFS-2 and prepare for the newly created VMware LUN.

   Note: You can also delete the partition by clicking **Edit…** and **Remove…**, depending on the current state of the existing LUN. Repeat for every LUN of the external virtual HBA, except for the local LUN of internal virtual HBA.

   Verification: Refresh if necessary once all existing LUNs are removed. Under the same **Disks and LUNs** submenu, verify each external LUN is in the raw format and displays the following message on the far right of the LUN ID:

   Create Volume…

3. Click **Create Volume…** (see Figure 4.3).

4. Select **Typical** to automatically partition and initialize the entire VMFS-2 disk partition. By choosing **Typical**, VMware automatically partitions the whole LUN for the created VMFS-2 volume (see Figure 4.4). Otherwise, click **Custom** to customize partitioning.

5. Optional: Enter a name (such as My_Lun01) for the LUN label.

6. From **Volume Type**, select the file system VMFS-2, VMware Core Dump or Extended Partition (default is VMFS-2).
7. Click **Extended Partition** if you plan to create more than four VMFS-2 volumes from the LUN.
8. Enter the appropriate volume label (see Figure 4.5).
9. Enter a size for the volume capacity (see Figure 4.6). Default option is recommended for file size. Otherwise, choose the appropriate maximum file size as the storage LUN capacity is available. Maximum file size is 28.50 terabytes. Public access option is fine and share access is used for clustering.
10. Click **Finish** to create the volume.
11. Once the task is completed and the window refreshed, verify that the volume has been created.
12. Verify from **Disks and LUNs** menu, under each LUN, appropriate LUN ID with LUN label and VMFS-2 file system.

My_Lun01 (VMFS-2.11 Volume)	Edit…

For raw device LUN which does not have VMFS-2 file system, should only appear as

Free Space	Create Volume…

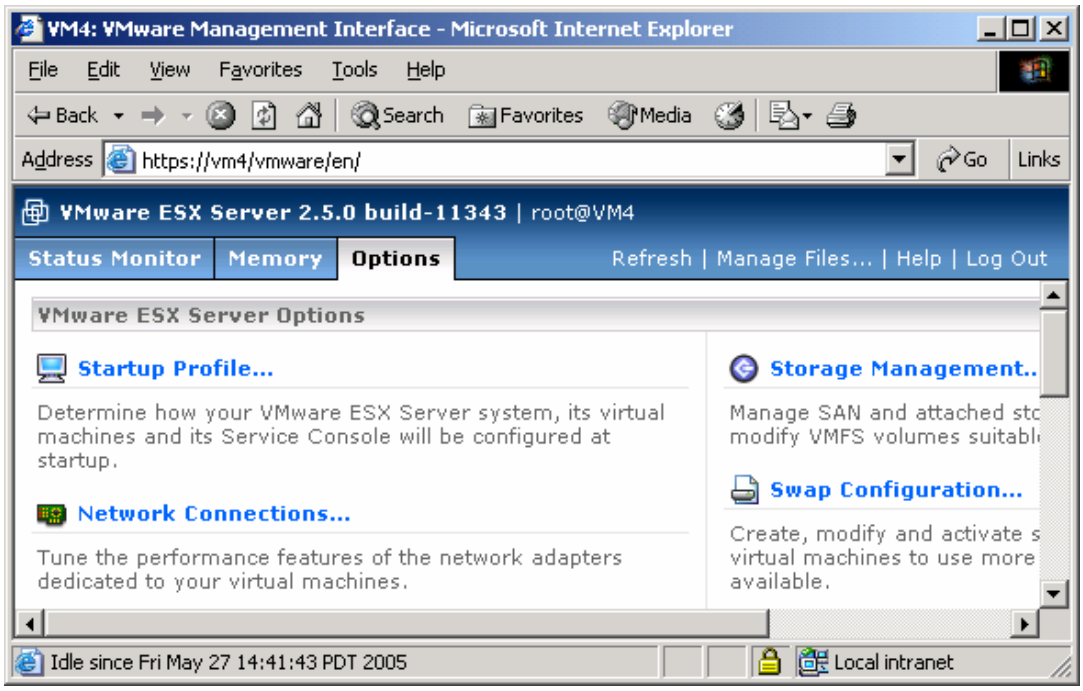

Figure 4.2 Loading the VMware Storage Management Options Module

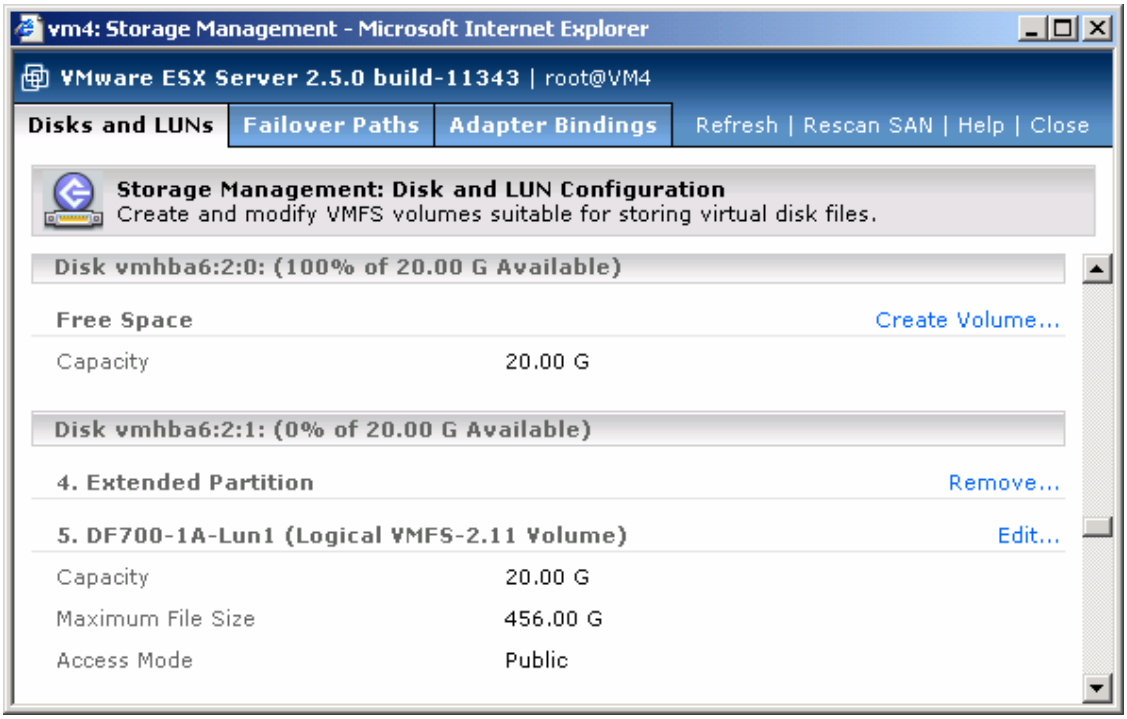

Figure 4.3 Accessing VMware Disks and LUN Management

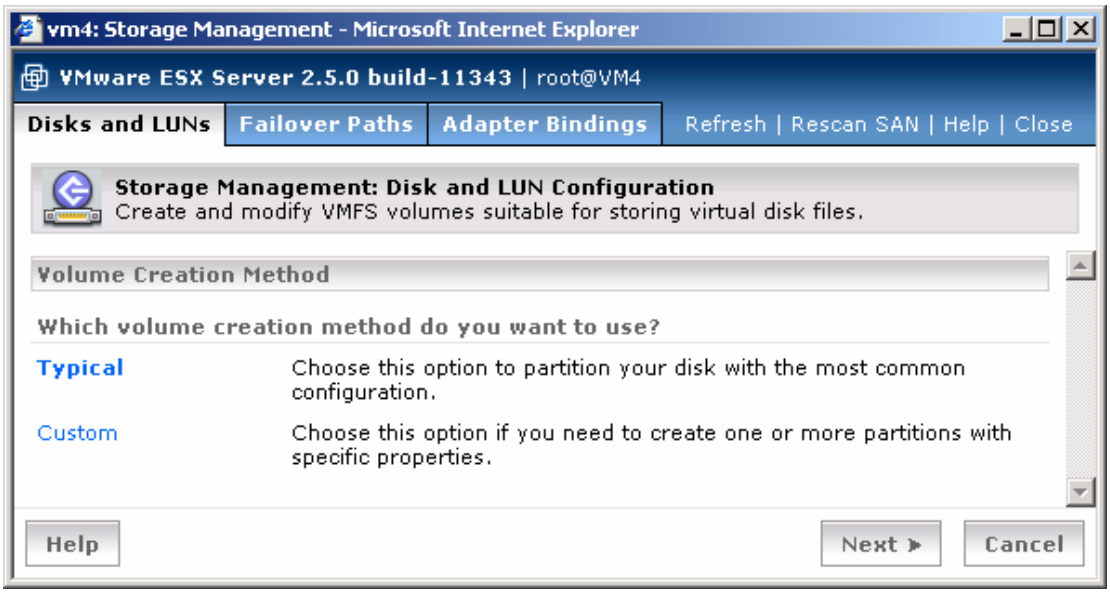

Figure 4.4 Customizing Disk and LUN Configuration

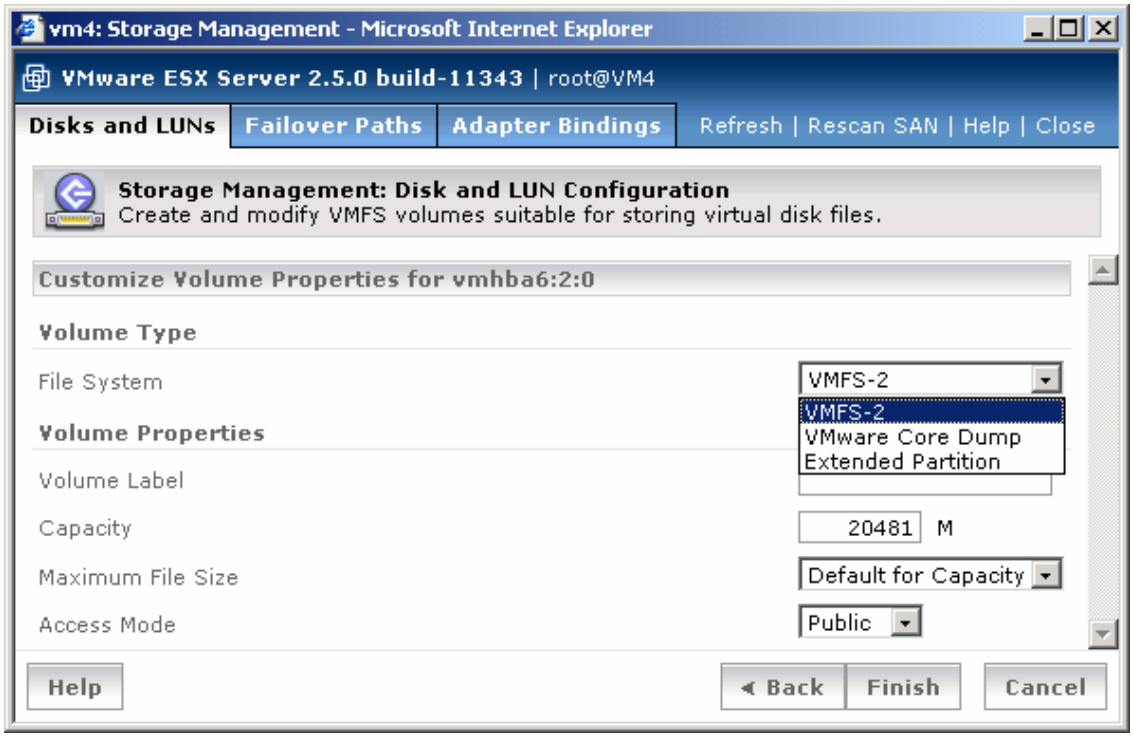

Figure 4.5 Selecting a File System for Volume Creation

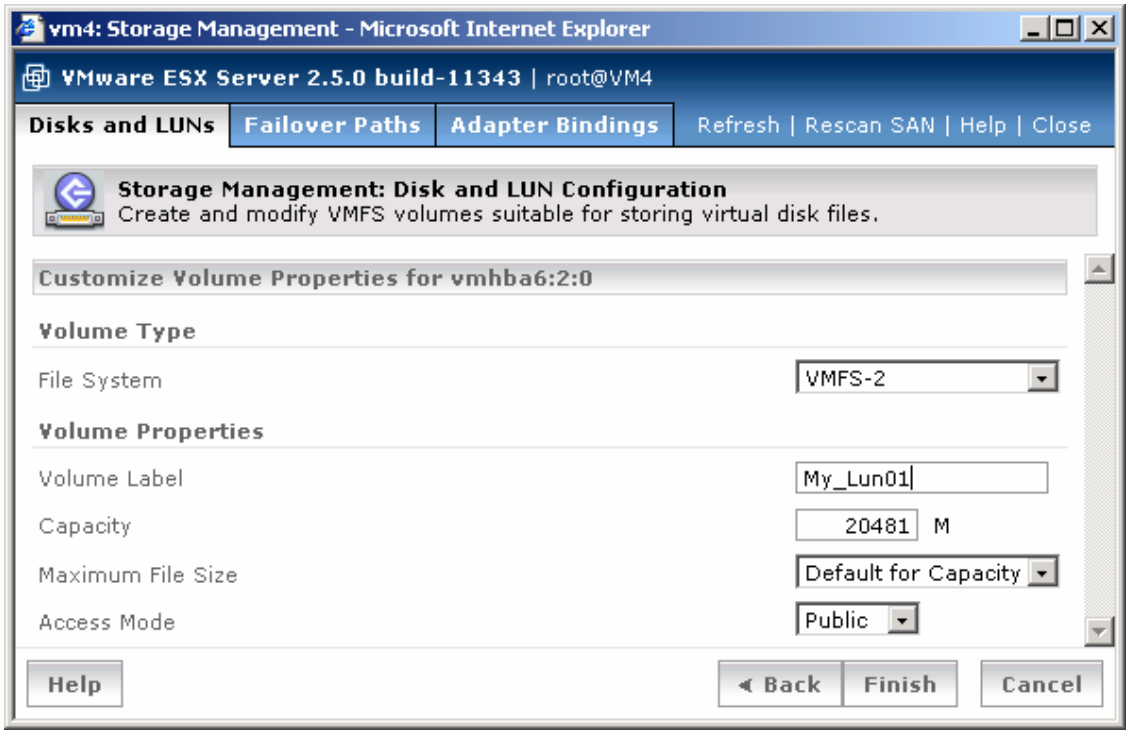

Figure 4.6 Creating the VMFS-2 Volume

4.3 Assigning VMFS-2 and Raw Devices to the Guest OS

After you create the VMFS-2 volumes, assign the new VMFS-2 volumes and raw devices to Guest OS, so the VMware Guest OS can discover and manage the new devices.

To assign new VMFS-2 and raw devices to the Guest OS:

1. Close the **Storage Management** and open the **Status Monitor** menu to assign created VMware volumes to a virtual machine (see Figure 4.7).
2. Under the **Virtual Machine** section, verify that the virtual machine is power off (a red square icon should appear before the LUN is assigned to it).
3. Click the down arrow next to the power button and click **Power Off** to turn off the virtual machine or **Configure Hardware…** to start assigning LUN(s) to virtual machine (see Figure 4.8).
4. Under **Hardware** menu, click on **Add Device…** and select **Hard Disk** option (see Figure 4.10). The **Virtual Disk** window appears, with three options for adding a LUN (see Figure 4.11):

   i. Blank

   ii. Existing

   iii. System LUN/disk
5. Select **Blank** to create a virtual disk with the image file extension .vmdk (see Figure 4.11).
6. Under **Image File Location**, select a VMFS-2 Volume (i.e., My_Lun01) and assign a VMware virtual disk image a name with a .vmdk extension (i.e., My_Virtual_Disk.vmdk.) Otherwise, "untitled" will be the assigned (see Figure 4.12).
7. Specify a size for the virtual disk image or leave the number as a default for the whole VMFS-2 volume to be partitioned as the virtual disk image (see Figure 4.13).
8. Select a Virtual SCSI Node for the newly added virtual disk image vmdk (see Figure 4.13). Use the form SCSI x:y, where x is the virtual SCSI node and y is the virtual disk image number. We recommend you use the default virtual SCSI node, as VMware sequentially increments the value for the next virtual disk image added to the VMware Guest OS.
9. For **Disk Mode** option, **Persistent** is the default. Changes to the virtual disk image are immediately and permanently written to the disk. Other options are **Nonpersistent**, **Undoable**, and **Append**.
10. Select **Blank** to create new virtual disk image from any VMFS-2 volume that has sufficient capacity (see Figure 4.15).
11. Verify that virtual disk is created and assigned as, under **Hardware** menu of the Guest OS virtual machine:

Virtual Disk	(SCSI x:y)
Device	VMware™ Disk Image
Location	My_Lun01:Virtual_Disk01.vmdk
Mode	Persistent

12. Click **System LUN/disk** to assign a raw device which has no VMFS-2 file system, to the virtual machine. A system/LUN disk or raw device will be in the form vmhbax:y:z, where:

    X : vmhba number

    Y : target SCSI ID

    Z : LUN number

    There are two ways to select a raw device:

    – As a physical device for assigning to a Guest OS. VMware presents to the Guest OS the raw device just like a raw LUN from a storage system.

    – As a disk meta file. A disk meta file acts as a virtual disk image from a VMFS-2 volume that helps VMware track all raw device target SCSI ID while allowing the Guest OS to access the underlying hardware directly. This represents VMware's physical compatibility mode that tracks all raw device target SCSI IDs and allows the Guest OS to take advantage of disk mode and other features of virtual disk. This is VMware's raw device virtual compatibility mode.

    To select a raw device without using meta data file, select the desired target LUN/disk and click **OK**.

    To select a raw device with the usage of meta data file, click the **Meta Data** box, select a VMFS-2 file for a meta data file location, and create a virtual disk image as a meta data file with an extension vmdk. Select a virtual SCSI node and a physical or virtual compatibility mode with appropriate disk mode for the raw device. Click **OK** to create a VMFS-2 managed raw device to assign to a VMware Guest OS.

    Verification: Verify that raw device is created and assigned as, under **Hardware** menu of the Guest OS virtual machine:

Virtual Disk	(SCSI x:y)
Device	System LUN/Disk
Location	My_Meta_Data_Disk01:Meta_Data_File01.vmdk
Mode	Physical

13. Click **Select Existing** to select an existing unassigned .vmdk disk image (i.e., My_Virtual_Disk) to a virtual machine. Since the virtual disk image has been created, clicking **Select Existing** may prevent you from being able to size the virtual disk image.

14. Close the **Hardware Configuration** window to complete the creation of the virtual disk and its assignment to the Guest OS virtual machine.

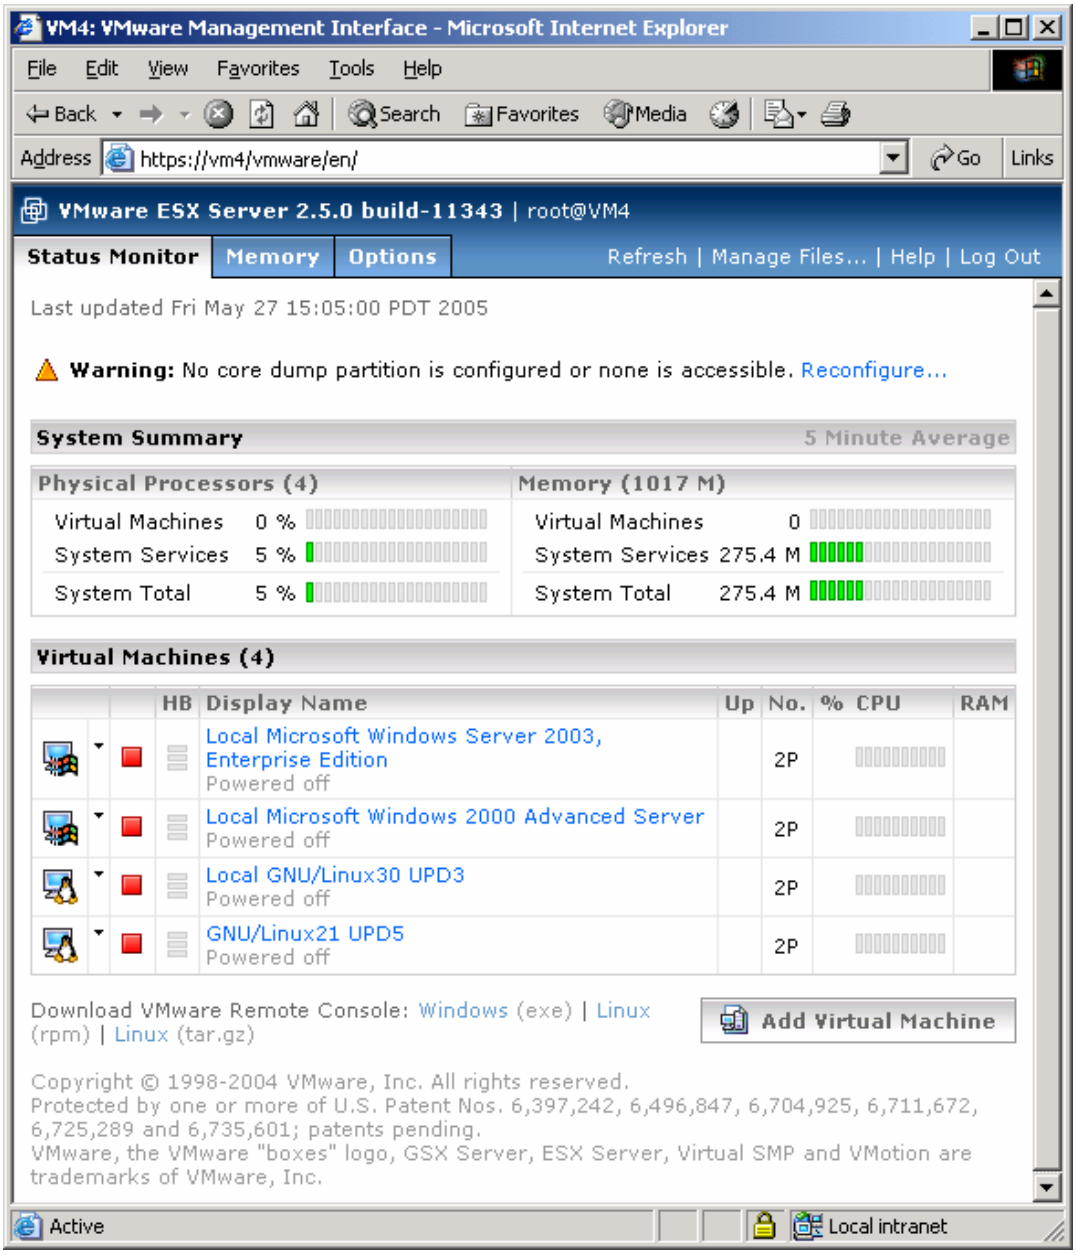

Figure 4.7 VMware Guest OS Status Monitor

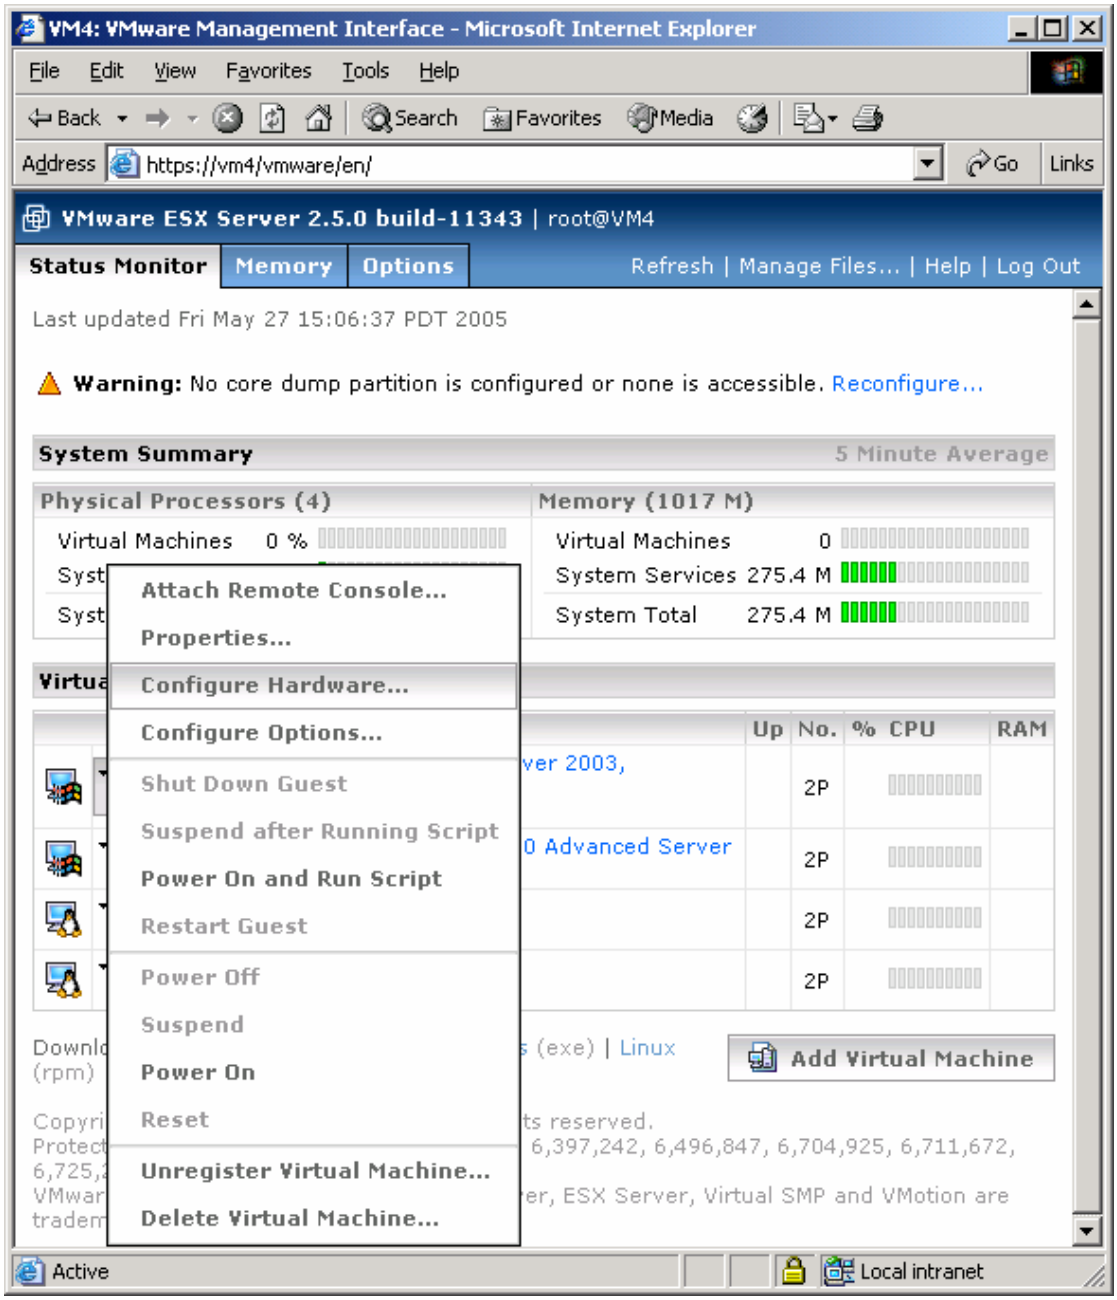

Figure 4.8 Guest OS Hardware Configuration

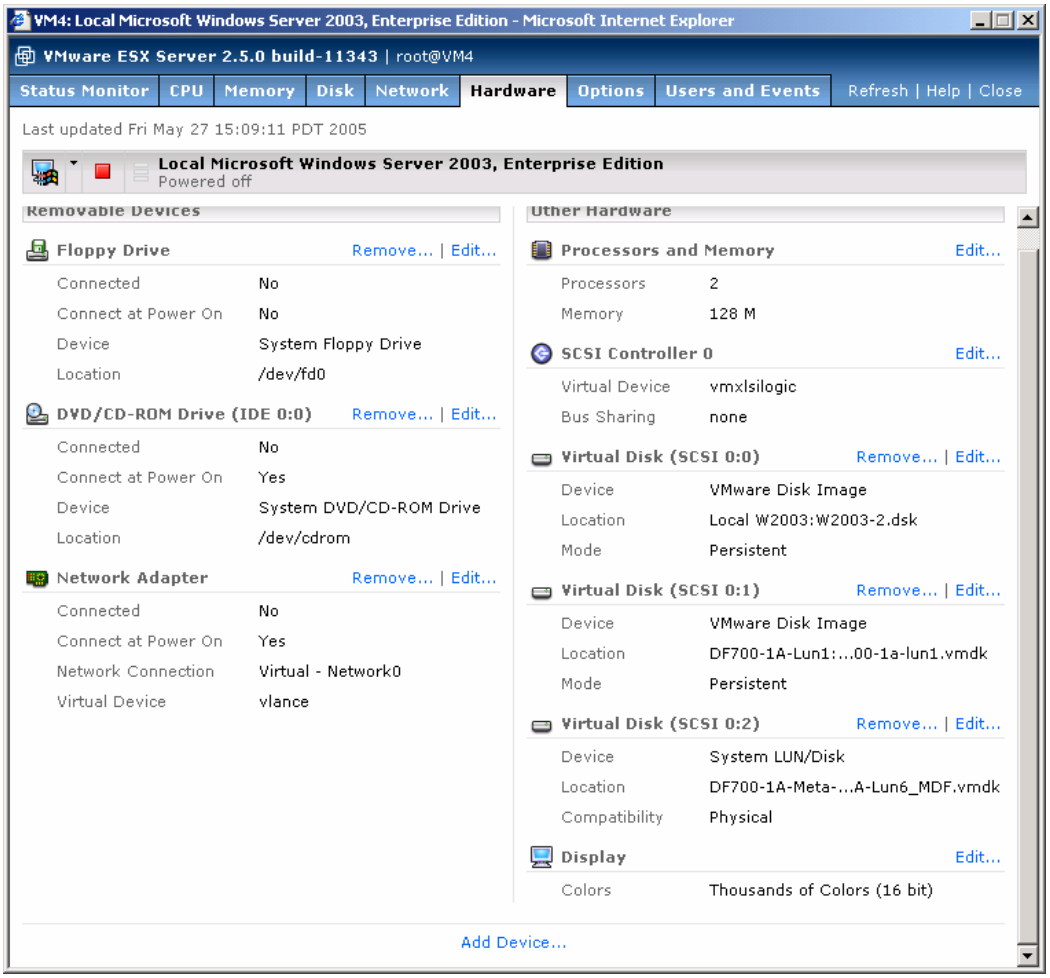

Figure 4.9 Guest OS Device Addition

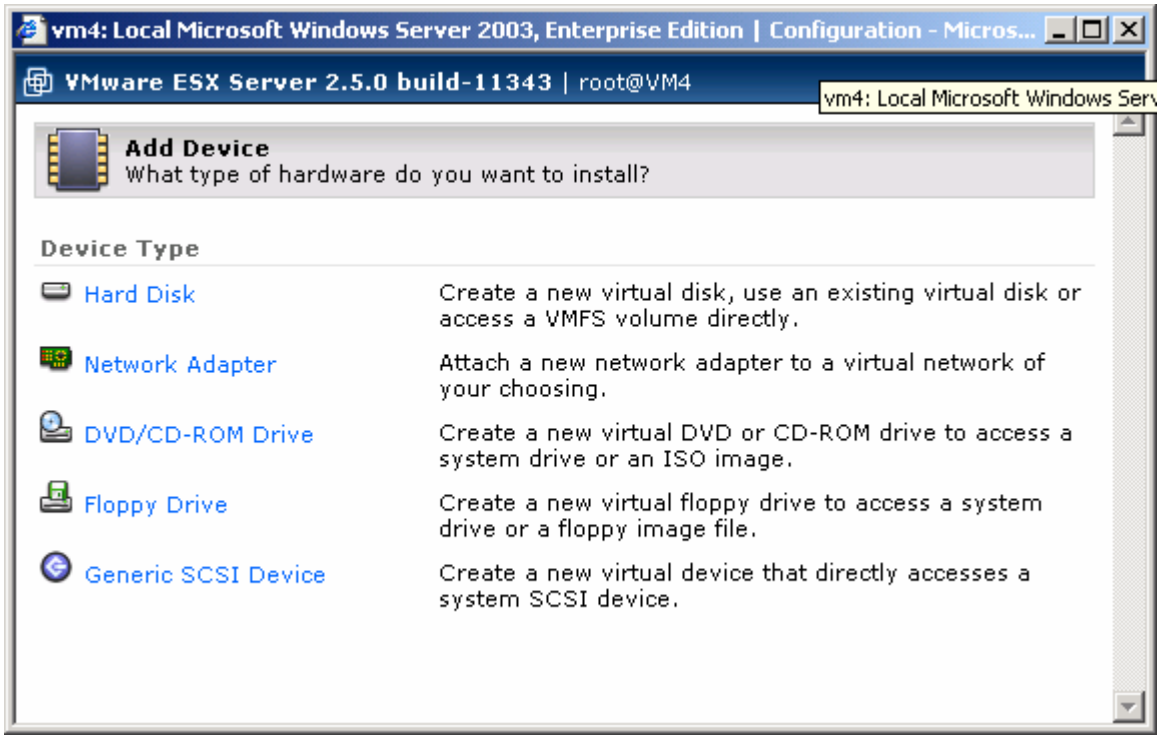

Figure 4.10 Guest OS Device Type Selection

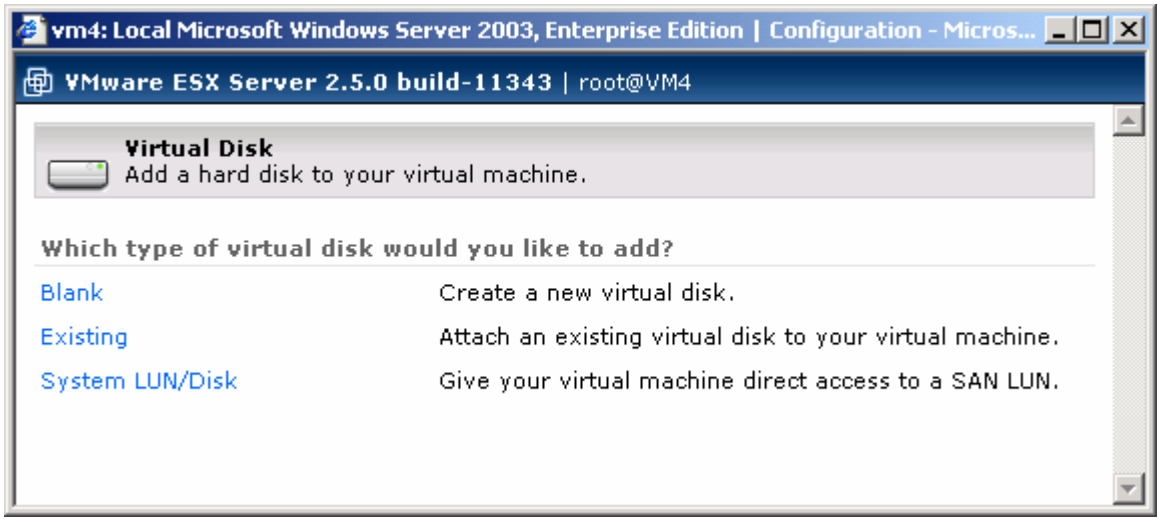

Figure 4.11 Guest OS New Virtual Disk Type Selection

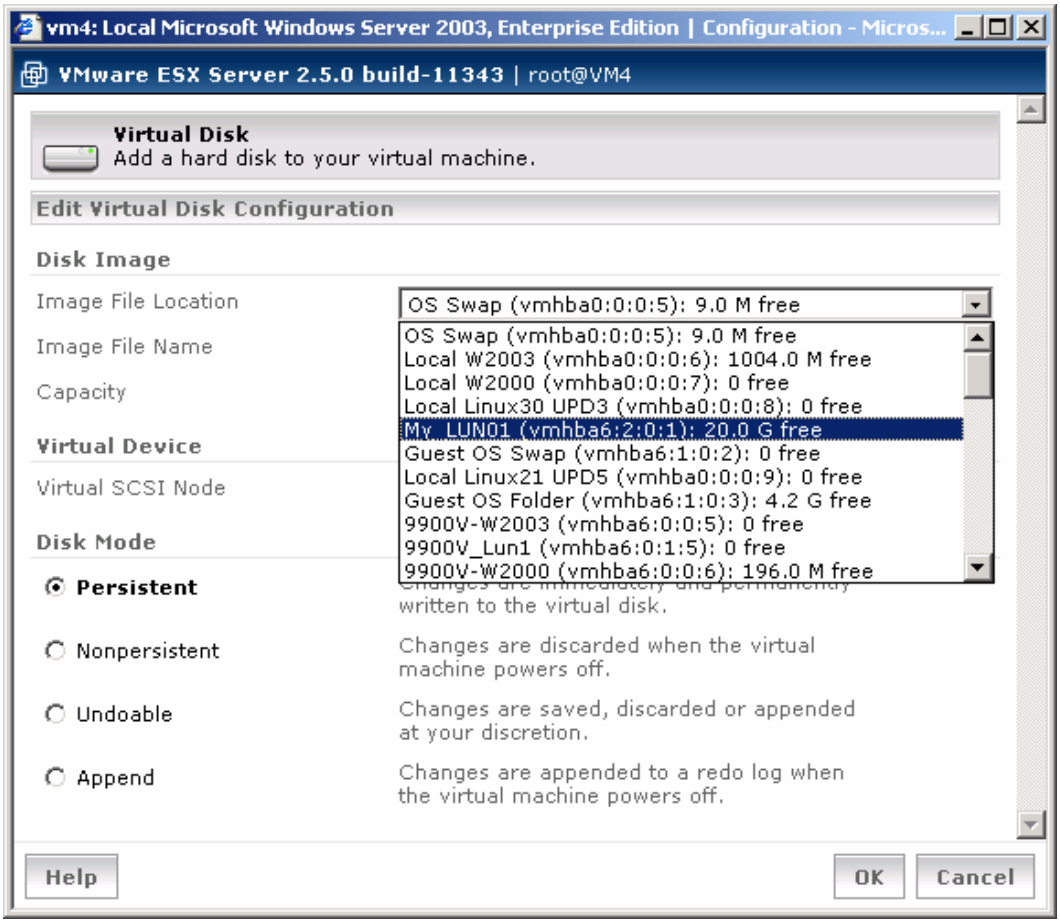

Figure 4.12 Guest OS New Virtual Disk Selection

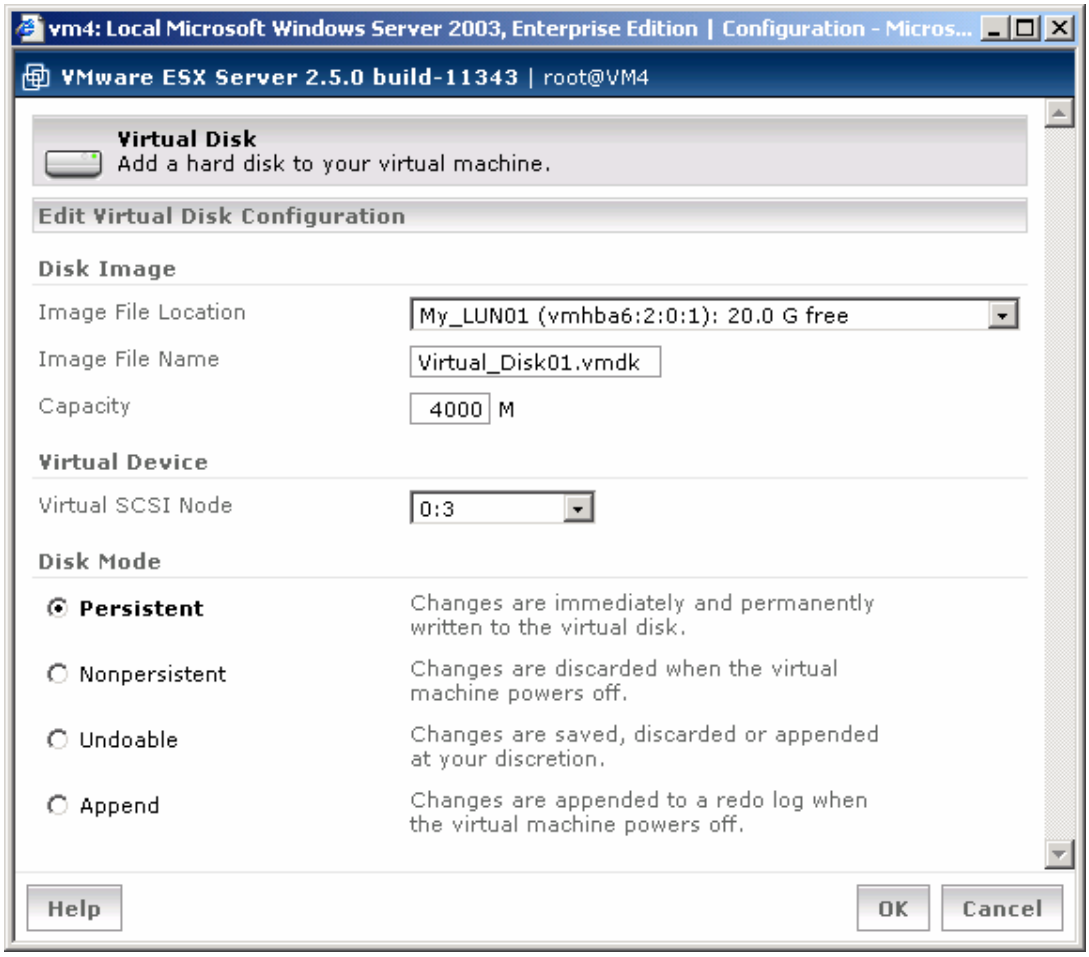

Figure 4.13 Guest OS Virtual Disk Sizing

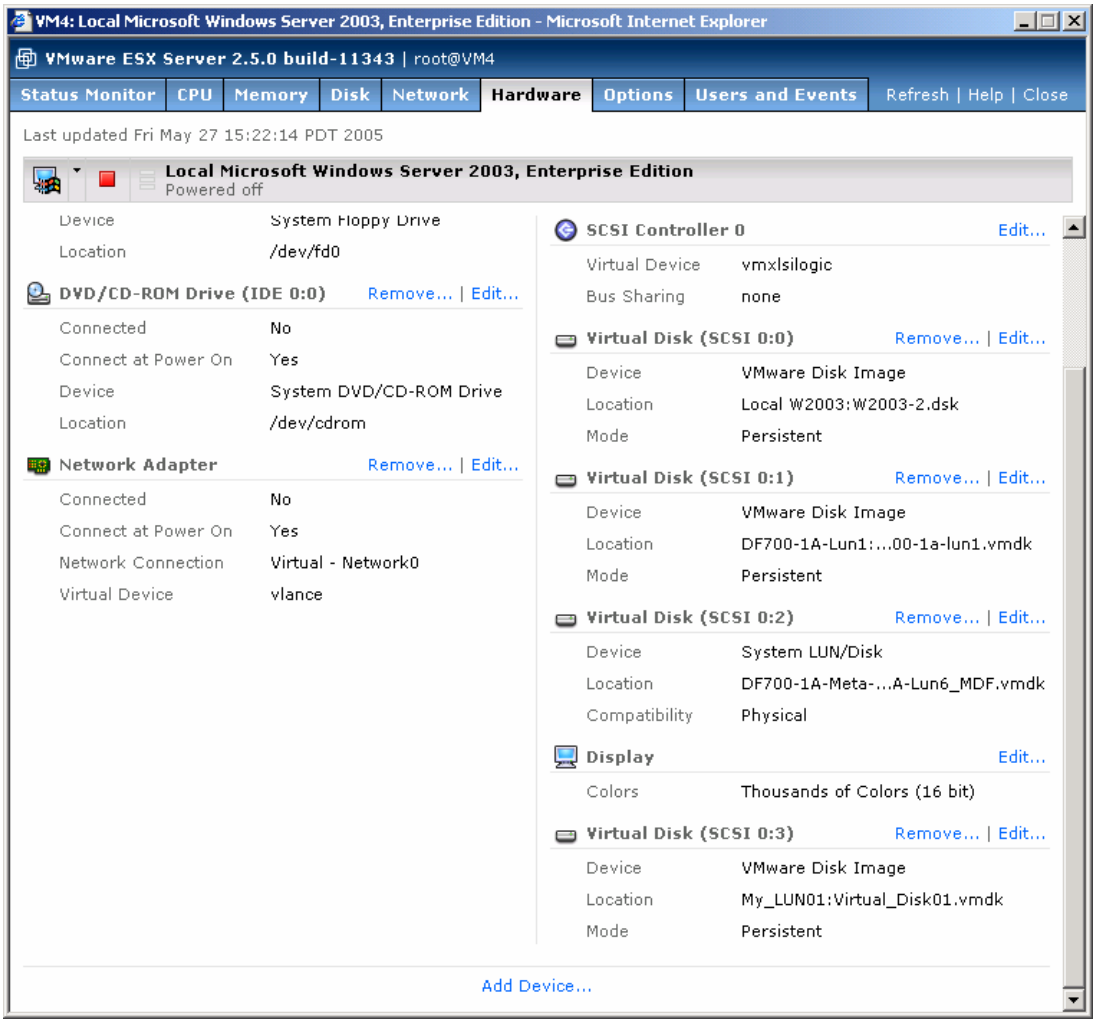

Figure 4.14 Guest OS New Virtual Disk Addition Completion

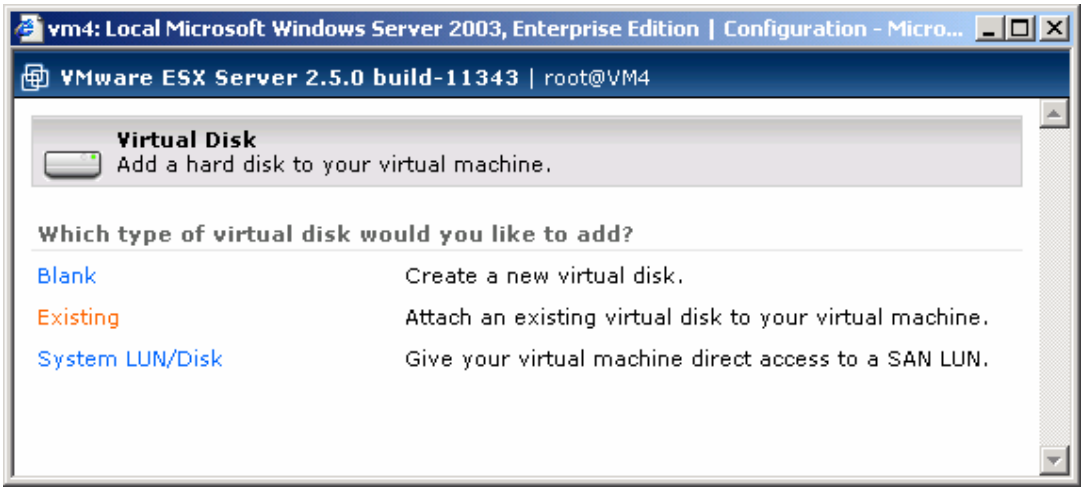

Figure 4.15 Guest OS Existing Virtual Disk Type Selection

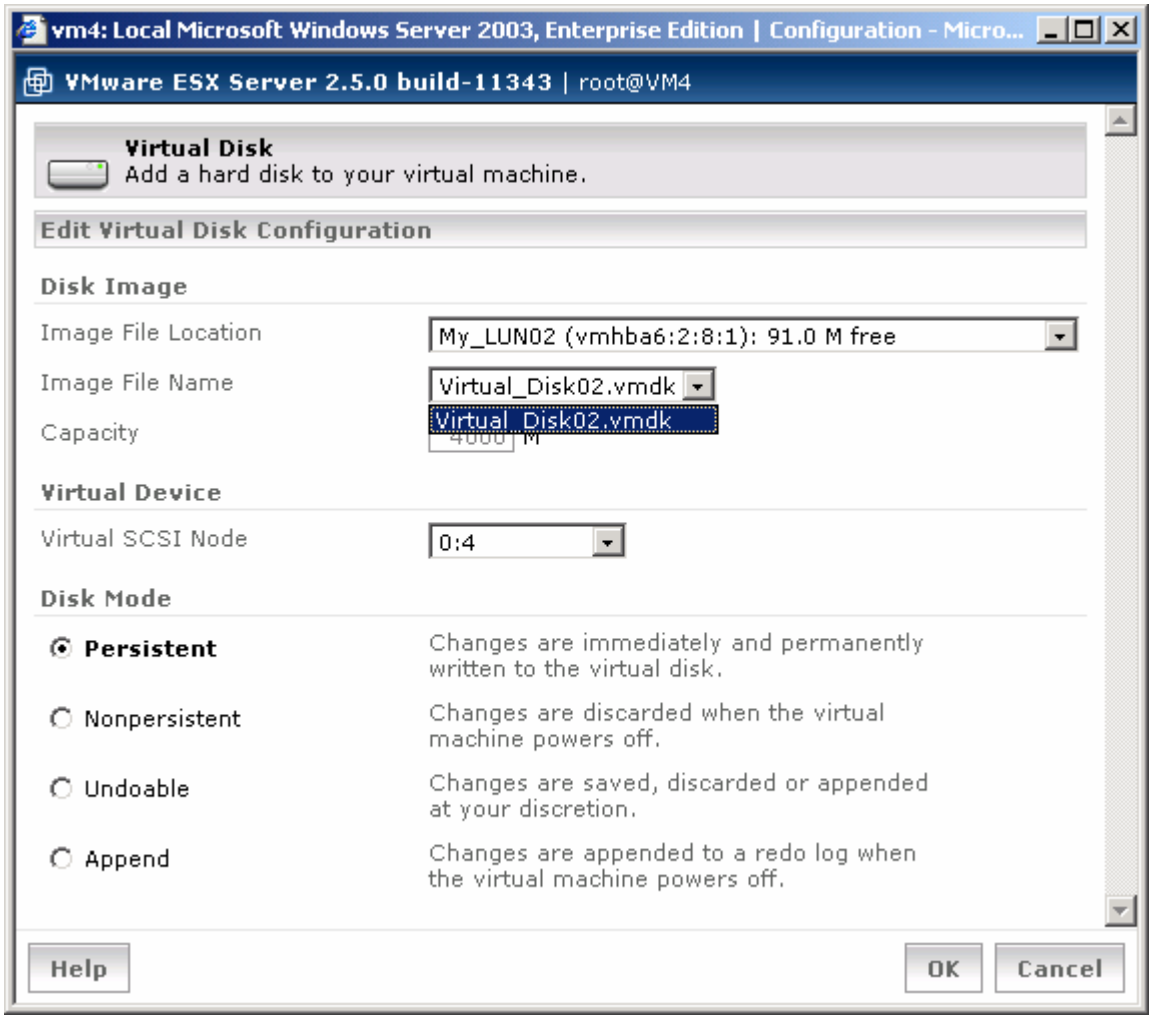

Figure 4.16 Guest OS Existing Virtual Disk Selection

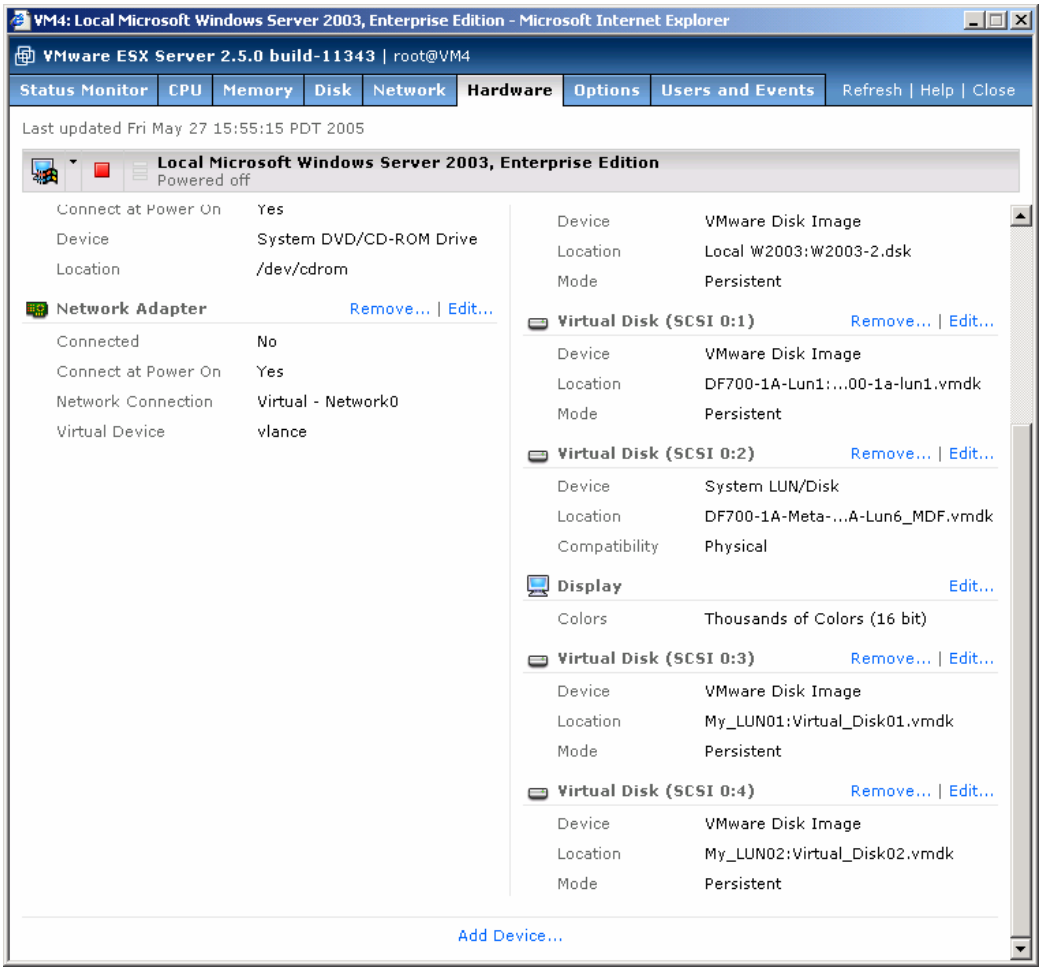

Figure 4.17 Guest OS Existing Virtual Disk Creation Completion

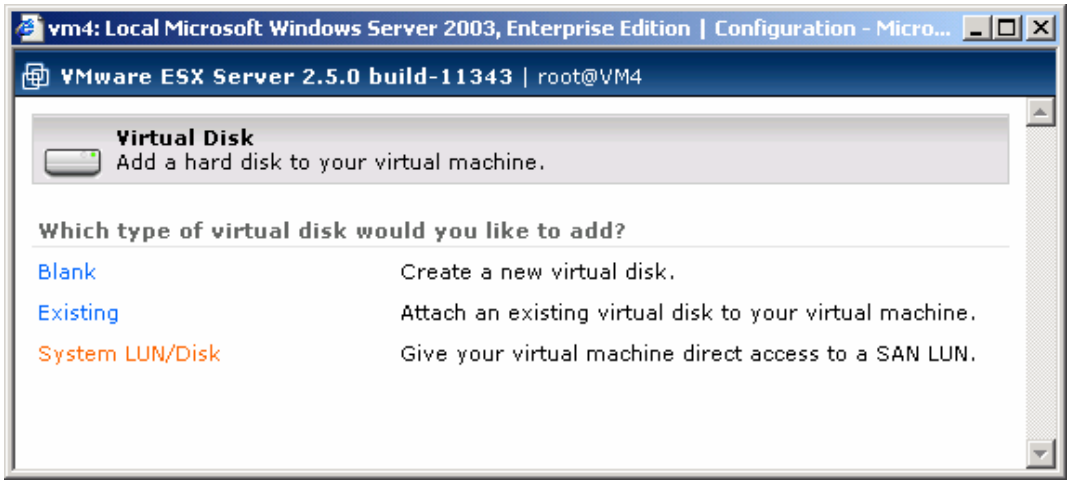

Figure 4.18 Guest OS Raw Device Virtual Disk Type Selection

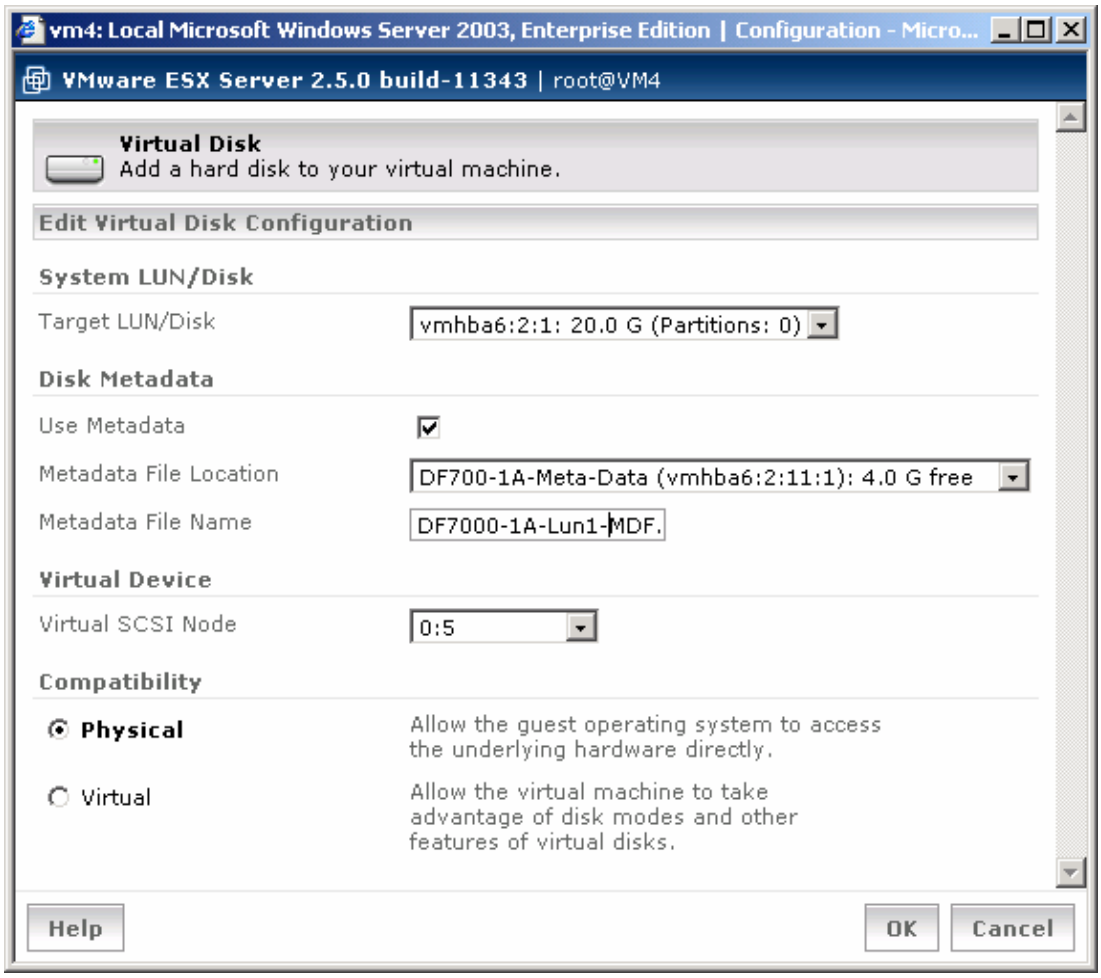

Figure 4.19 Guest OS Raw Device and Meta Data File Selection

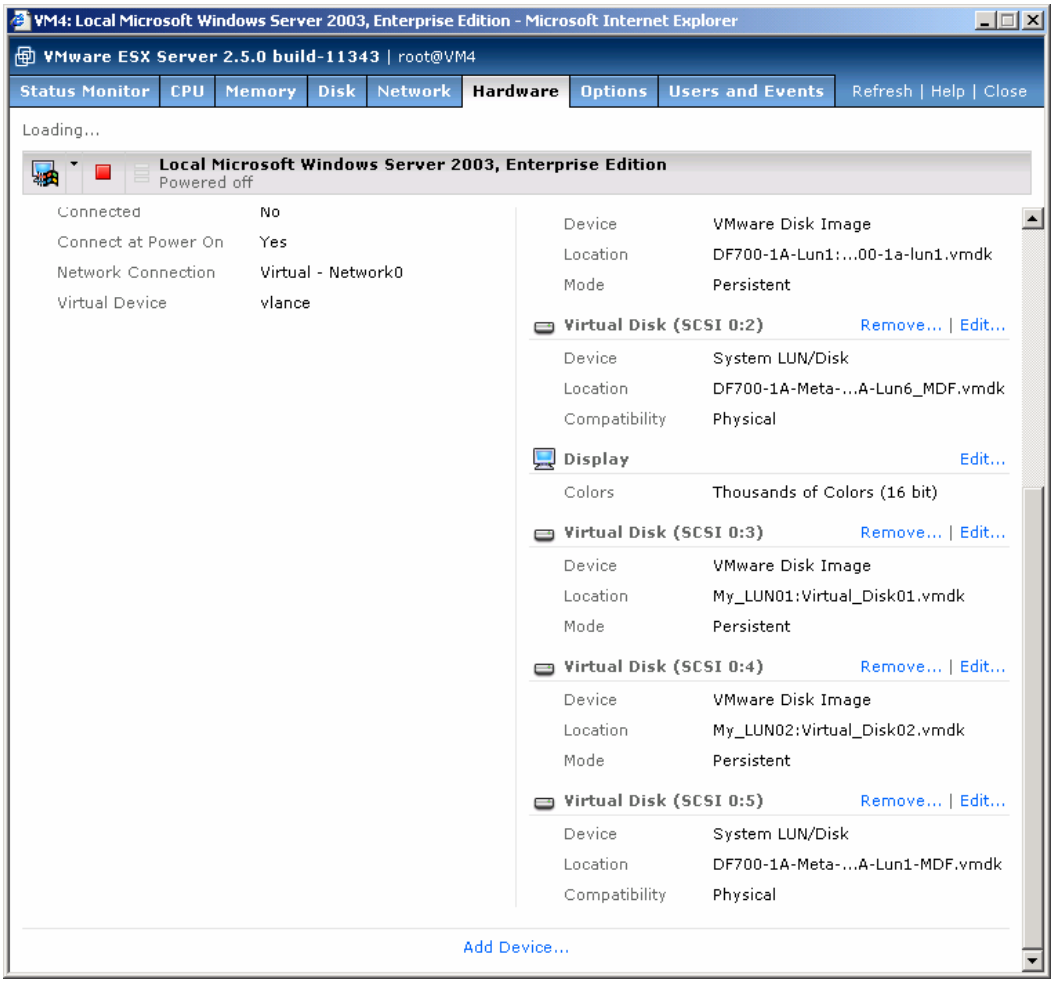

Figure 4.20 Guest OS Raw Device Addition Completion

4.4 Guest OS Management

After you assign the virtual disks to VMware Guest OS, the Guest OS can be powered on and start discovering and managing the new devices. The Guest OS can be managed with the following GUI configuration options under Guest OS properties:

- Status Monitor
- CPU
- Memory
- Disk
- Network
- Hardware
- Options
- Users and Events

Most storage-related and device peripherals configuring options are found under the **Hardware** tab, while Options manages Guest OS display name and Guest OS virtual machine startup and shutdown options at ESX Server system startup and shutdown. Guest OS virtual memory allocation can also be adjusted under the **Hardware** menu (see Figure 4.22).

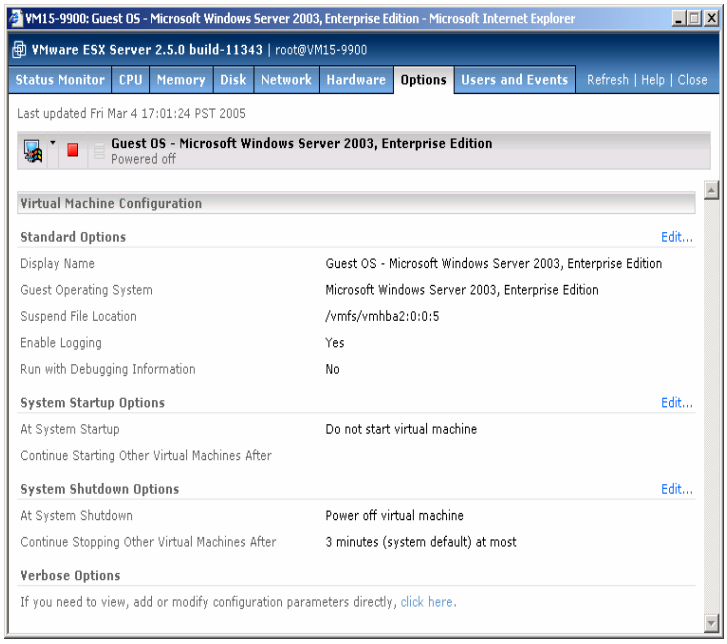

Figure 4.21 Guest OS Virtual Machine Options Configuration

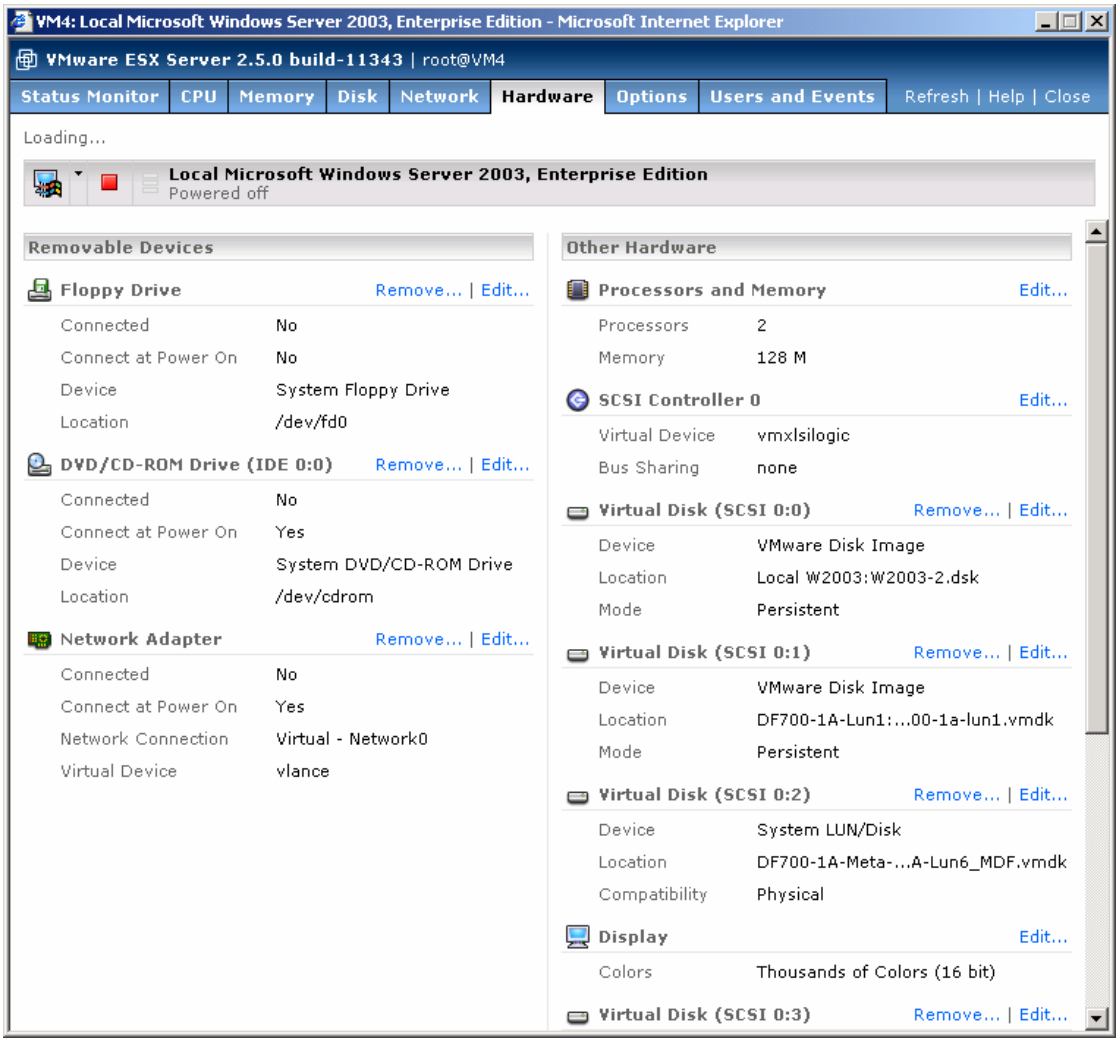

Figure 4.22 Guest OS Virtual Machine Hardware Configuration Interface

4.5 Verifying Guest OS Access

The last step in new device configuration is to verify that the VMware™ Guest OS can access the new virtual disk images. To verify Guest OS access, power on Guest OS and start to discover LUN as usual as discover LUN with regular OS (i.e., Windows 2000 Advanced Server, and Red Hat Linux 3.0).

1. Click the red square button on the left on the Guest OS display or through **Power On** from drop-down menu of the remote console (see Figure 4.23, Figure 4.24, and Figure 4.25).
2. Download the appropriate VMware remote console, if necessary, and attach it to the powered-on Guest OS (see Figure 4.25 and Figure 4.26).

   Note: The Guest OS will not power on if there is an assigned virtual disk, but not attached to ESX Server (i.e., the virtual disk has been removed from the ESX Server or the virtual disk label has been altered.)
3. When the Guest OS starts, perform the necessary disk administration to discover the LUN.
4. For Windows operating systems, initialize and format the LUN through Disk Management. Once the LUN is successfully formatted, I/O can be executed normally.
5. Access the Disk Management and format LUN:
6. Verify that the LUN has been 100% successfully formatted with NTFS (see Figure 4.28).

   xx.xx GB NTFS

   Healthy
7. For Linux operating systems, follow these steps to discover a new LUN, make the file system, and mount the LUN once it is formatted (see Figure 4.28).
8. Verify that the Guest OS has discovered the LUNs:

   [...root]# dmesg |grep scsi | more
9. Verify that the output message indicates that the new device has been attached.

```
Attached scsi disk sda at scsi0, channel 0, id 0, lun 0
Attached scsi disk sdb at scsi0, channel 0, id 1, lun 0
…
```

   The message indicates that sda is the local internal disk of the Guest OS and sdb is the newly discovered device.
10. Perform a disk partition on device sdb:

    [... root]# fdisk /dev/sdb

    ...

11. Select the following options to create a partition for /dev/sdb:

    O : create a new empty DOS partition table

    n : add a new partition

    e : extended

    1 : partition number (1-4)

    1 : first cylinder

    <enter> : default for last cylinder

    w : write partition table

12. Make a file system for the device:

```
[…root]# mkfs –t ext3 /dev/sdb
```

13. Click **y** to make a file system for the entire device.

    Note: For Linux 2.1, execute the command `mkfs` with the option **–j** instead of **–t ext3** to create an ext3 file system

14. Make a directory and mount the file system:

```
[… root]# mkdir diskb
[… root]# mount /dev/sdb /diskb
```

15. Verify with `mount` command to verify /dev/sdb is mounted on /diskb as ext3 file system

```
[…root]# mount
…
/dev/sdb on disk /diskb type ext3 (rw)
```

16. Modify the fstab entry for the newly added LUN if necessary.

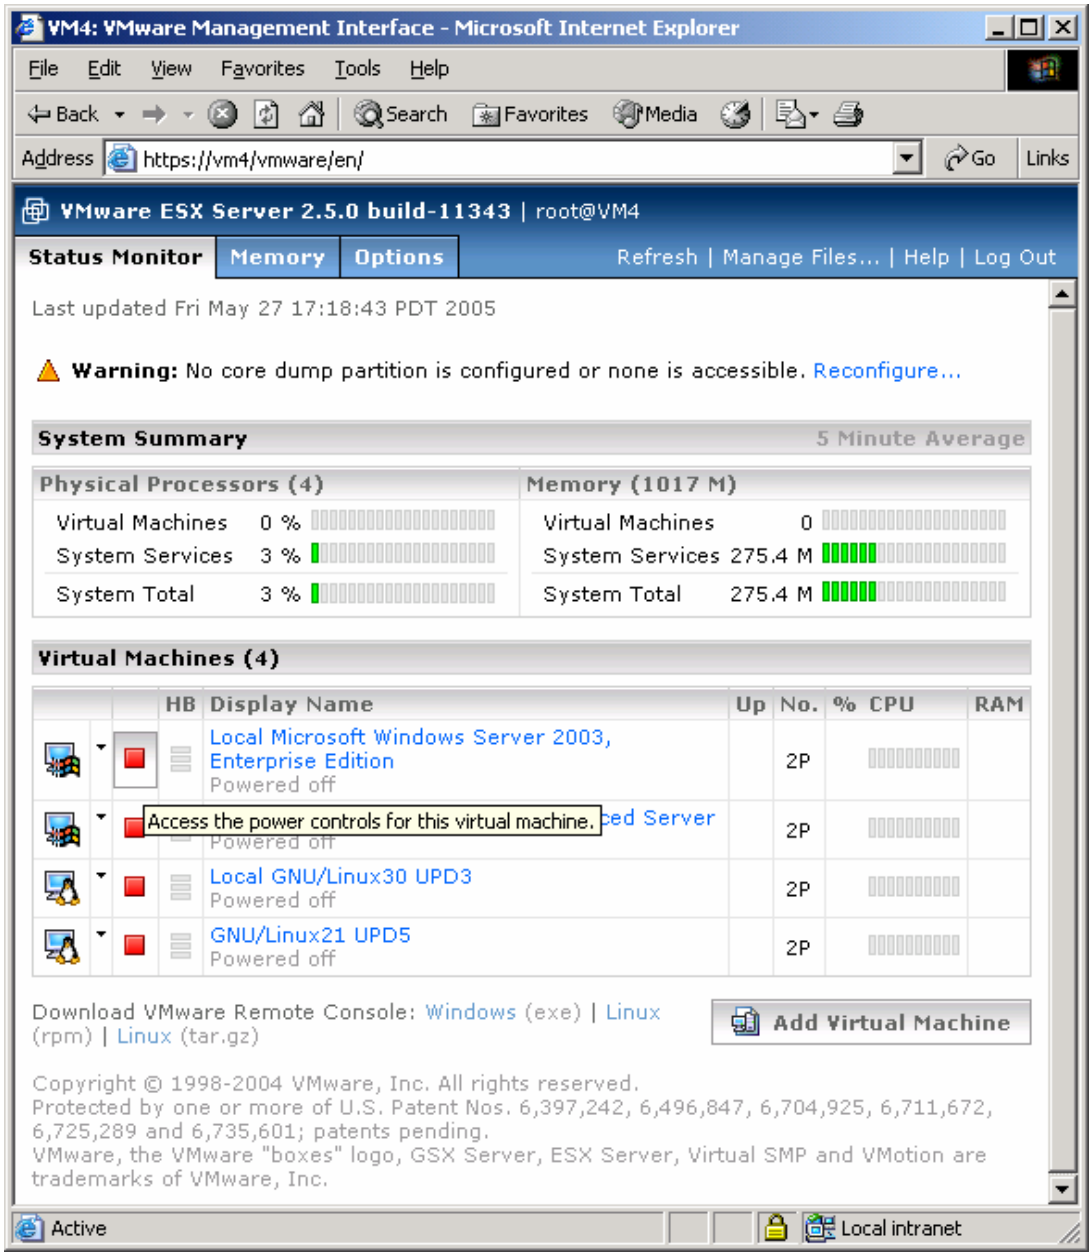

Figure 4.23 Access the Power Controls for Guest OS Virtual Machine

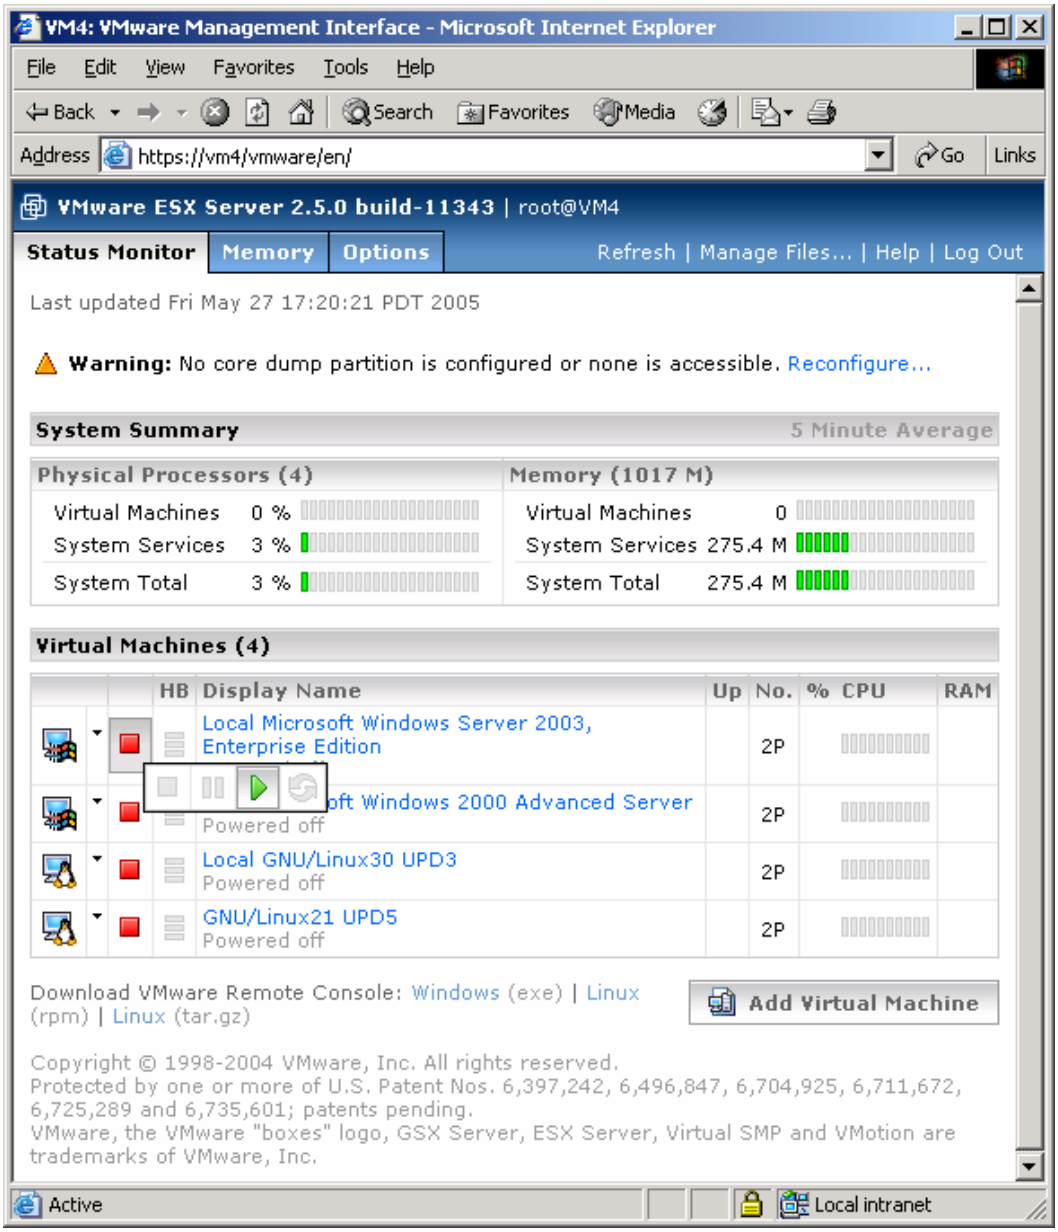

Figure 4.24 Guest OS Power On

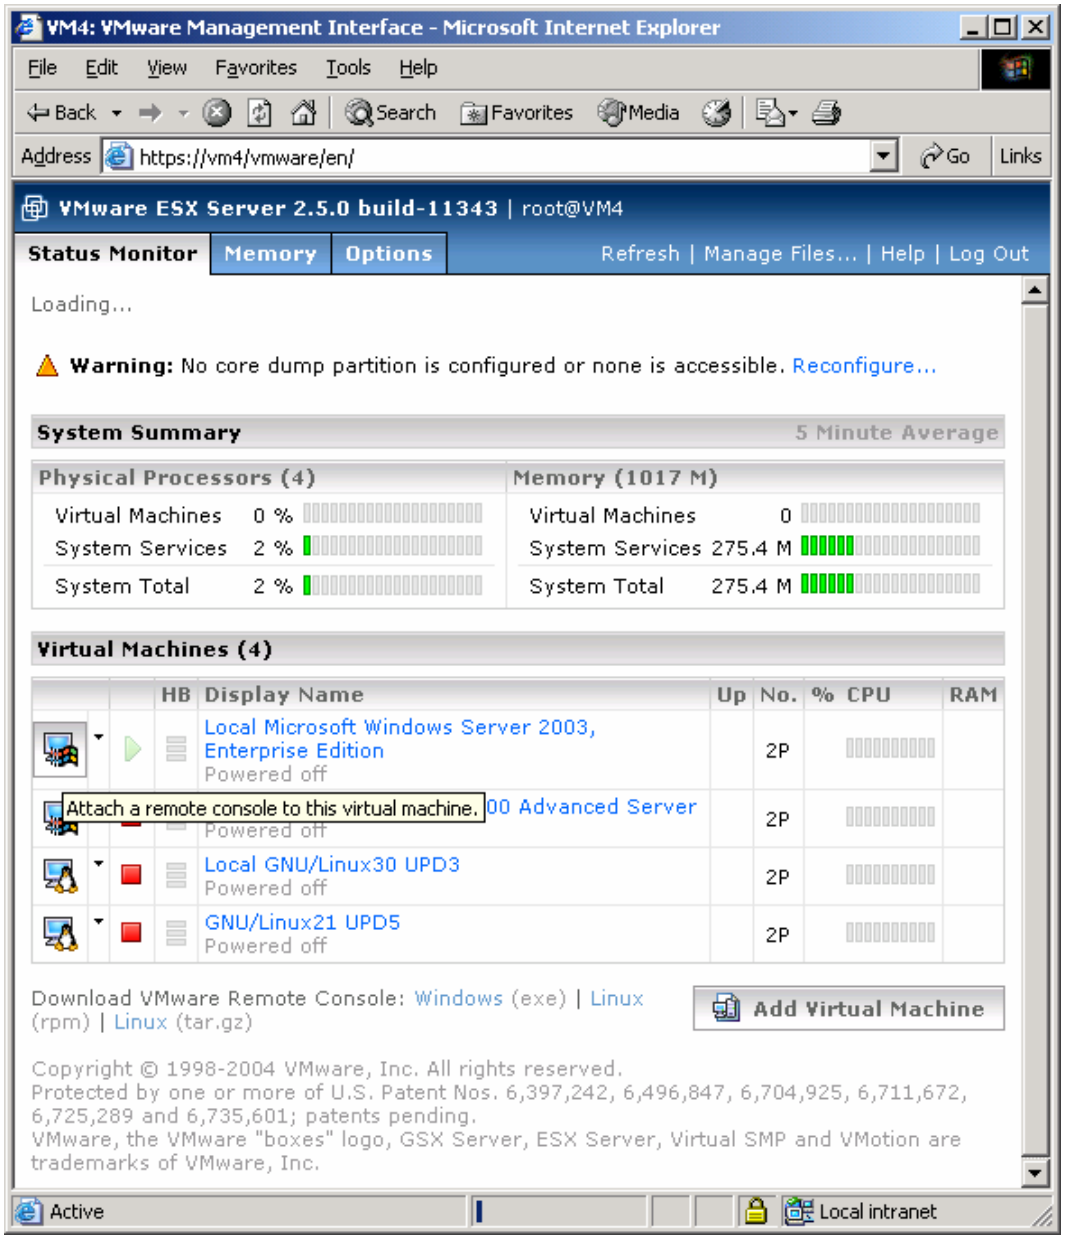

Figure 4.25 Attach a Remote Console to Guest OS

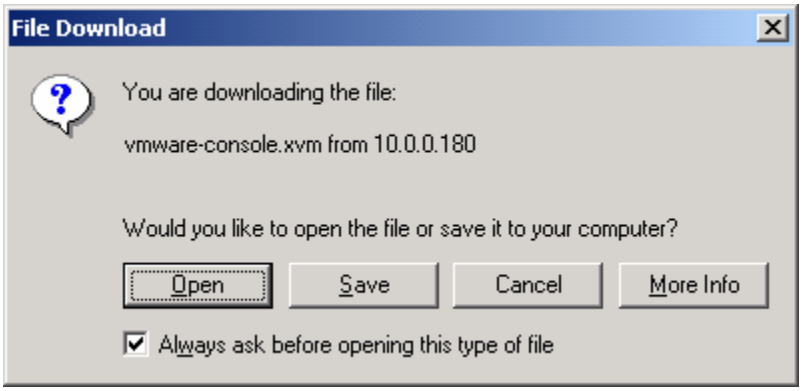

Figure 4.26 Open Remote Console

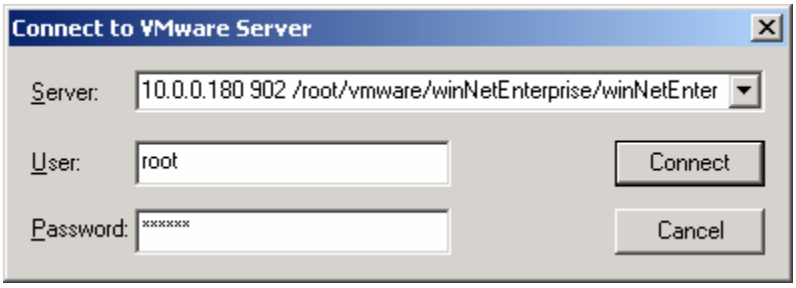

Figure 4.27 Connect Remote Console to Guest OS

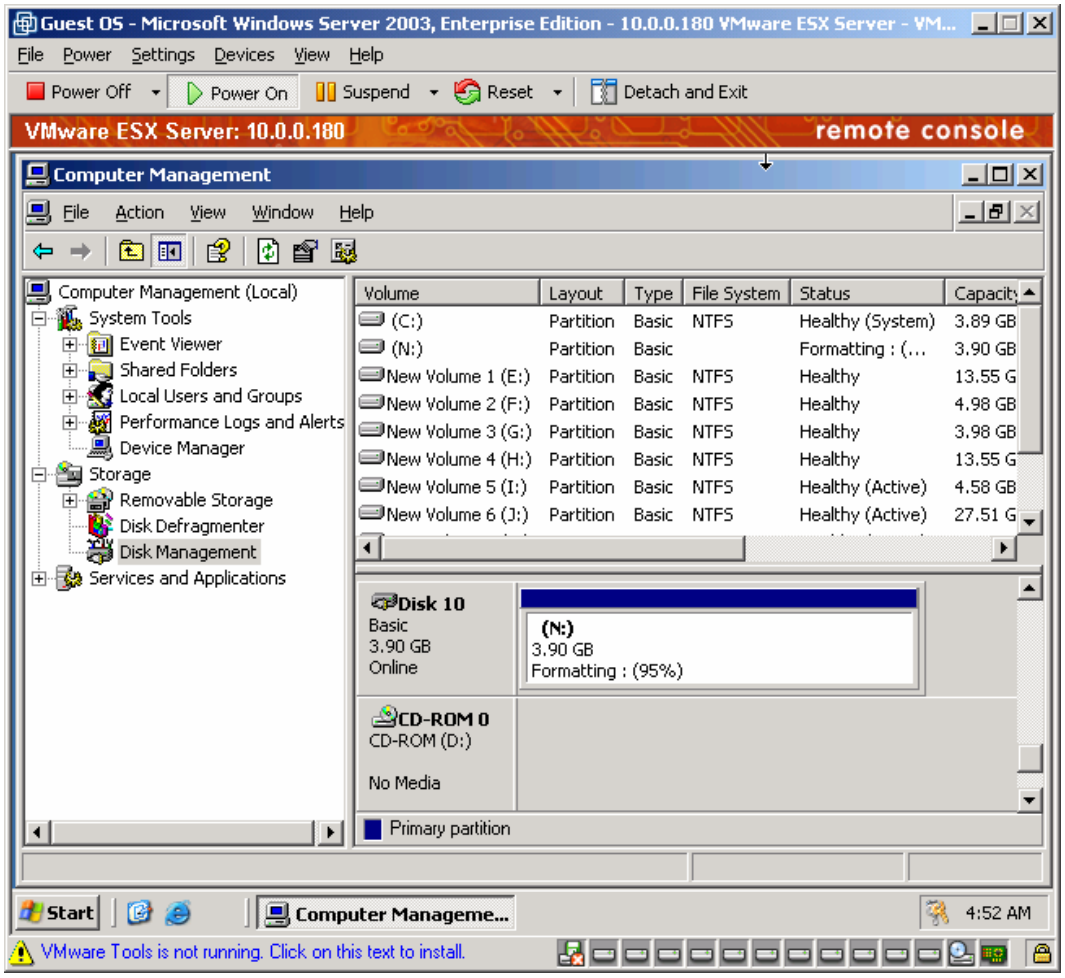

Figure 4.28 Windows NTFS Formatting

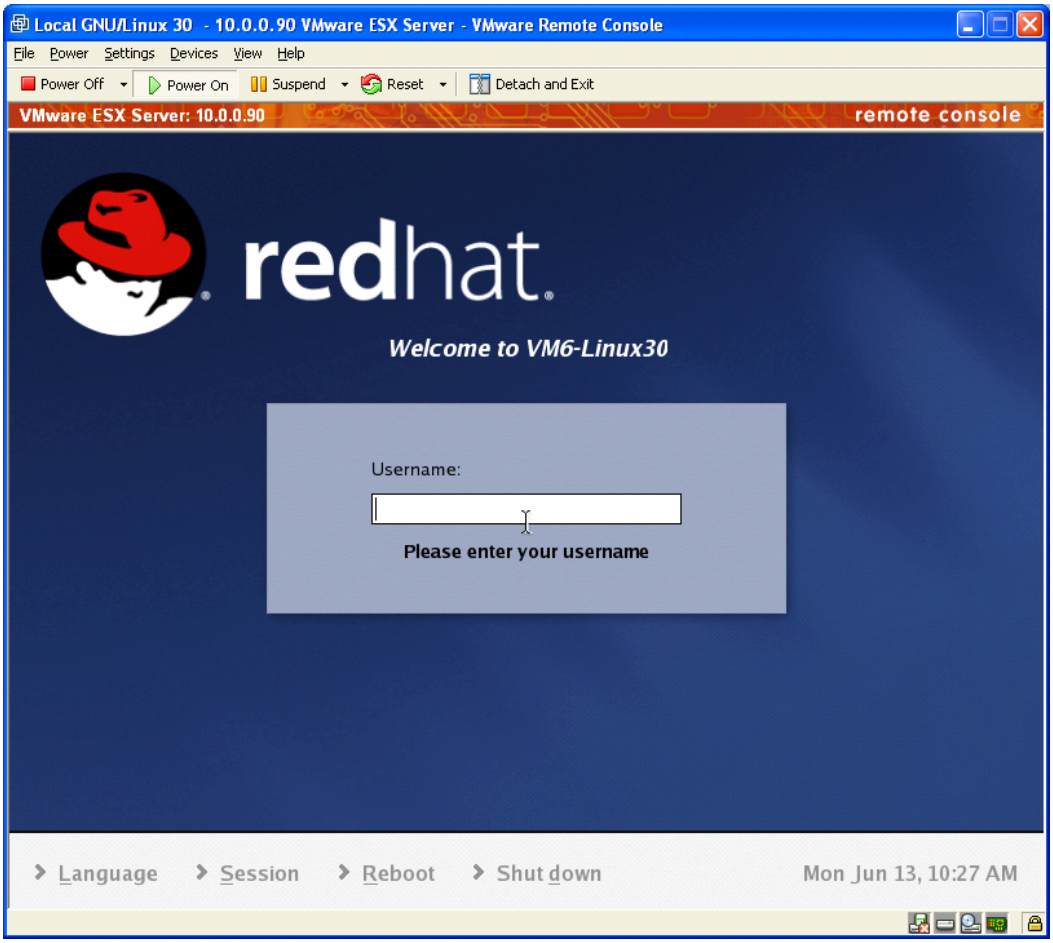

Figure 4.29 Red Hat Linux 3.0 Guest OS

4.6 Online LUN Addition/Deletion

VMware ESX Server supports online LUN addition and deletion for all Hitachi TagmaStore AMS/WMS systems. Consult with Hitachi Data Systems.

To add or delete LUNs to both paths:

1. Execute `ls` command to verify all LUNs on the primary and secondary paths.

```
[… root]# ls /proc/VMware/scsi/vmhba2
0:0  0:2  0:4  0:6  0:8
0:1  0:3  0:5  0:7  0:9
0:10
[… root]# ls /proc/VMware/scsi/vmhba3

stats

[… root]# cat /proc/VMware/scsi/vmhba2/0:0

Paths:fixed
     vmhba2:0:0 on*#
     vmhba3:0:0 on
```

Figure 4.30 LUNs Display

Execute the `vmkfstools –s` command at the ESX Server service console for each vmhba for ESX Server™ to scan for the newly added LUNs on both primary and secondary paths.

```
[… root]# vmkfstools –s vmhba2
[… root]# vmkfstools –s vmhba3
[… root]# cos-rescan.sh vmhba2
[… root]# cos-rescan.sh vmhba3
…
```

Figure 4.31 Using VMKFSTOOLS Command to Scan for Newly Added/Deleted Devices

2. The ESX Server should already discover the newly added/deleted LUNs. Executing the `ls` and `cat` commands may let you verify this. However, newly added/deleted LUNs may not be present to the VMware Management Interface until a rescan SAN function is performed:
   a) From the **Disks and LUNs** submenu of the Storage Management, click **rescan SAN**. Repeat the scan if necessary (see Figure 4.32).
   b) Display LUNs again at the service console to verify the newly added/deleted LUN.
3. From the **Disks and LUNs** tab or **Failover Paths** of the Storage Management, verify newly added LUNs (see Figure 4.33). At any time, you can execute the `ls` and `cat` commands to verify added/deleted LUNs. Repeat steps 1 through 3 if necessary.

Note: You can use the `vmkmultipath –q | more` command instead of `cat` to quickly display multiple LUN entries. Also, you can use the cursor key to repeat the previous command on the service console.

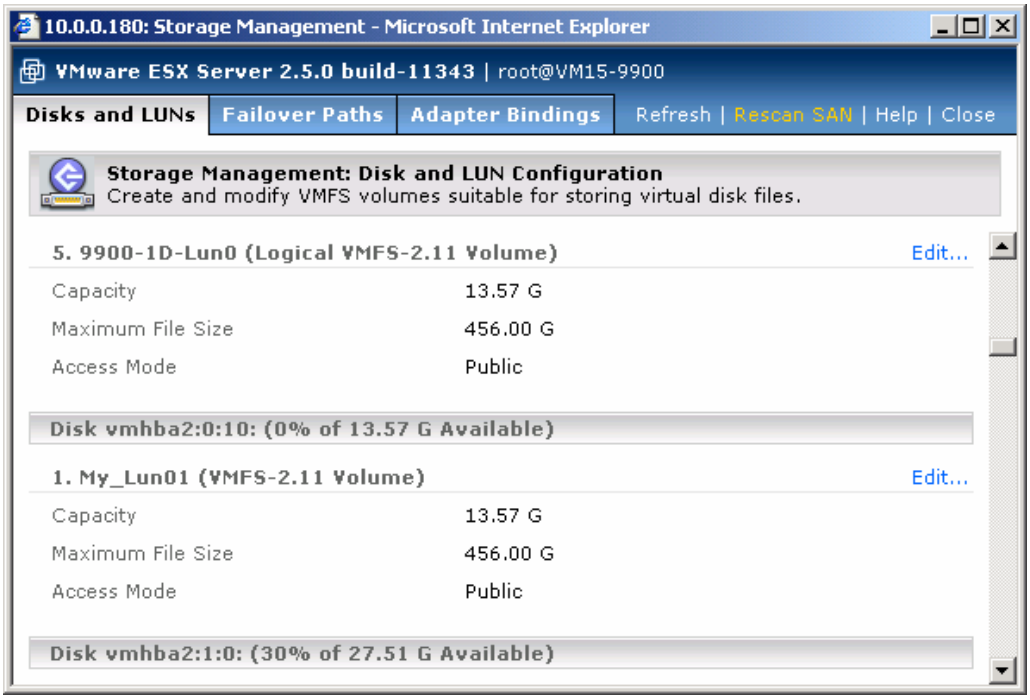

Figure 4.32 Rescan SAN

```
[… root]# ls /proc/VMware/scsi/vmhba2
0:0  0:2  0:6
0:1  0:3  0:7
0:10 0:4  0:8
0:11 0:5  0:9

[… root]# ls /proc/VMware/scsi/vmhba3

stats

[… root]# cat /proc/VMware/scsi/vmhba2/0:11

Paths:fixed
     vmhba2:0:11 on*#
     vmhba3:0:11 on
```

Figure 4.33 Discover Newly Added LUNs

Chapter 5 Failover and Fixed Policy

This chapter describes the failover and fixed policy. Topics include:

- Host Failover (section 5.1)
- Path Failover (section 5.2)
- Fixed Policy with Preferred Paths v.s. MRU (section 5.3)

Hitachi Thunder™ product supports industry-standard products and functions, which provide host and/or application failover.

For the VMware environment, the AMS/WMS supports the VMware Multi-Path Support (bundled with VMware ESX Server software product. Contact Hitachi Data Systems for more information.

5.1 Host Failover

Contact Hitachi Data Systems for the latest information about VMware supported host failover for the AMS/WMS and other Hitachi storage systems.

5.2 Path Failover

VMware's integrated Multi-Path Support is fully supported in VMware ESX Server 2.5. However, load balancing is **not** supported. To reestablish and verify the connection state of the primary path and secondary path after failover, verify the connection by executing `cat` or running the VMware multi-path display command at the service console. Repeat the command for each virtual device of the failing vmhba.

```
[… root]# cat /proc/VMware/scsi/vmhba2/0:0

Paths:fixed
     vmhba2:0:0 dead#
     vmhba3:0:0 on*


[… root]# cat /proc/VMware/scsi/vmhba2/0:1

Paths:fixed
     vmhba2:0:1 dead#
     vmhba3:0:1 on*
```

Figure 5.1 Path Failover Display

The output of the `cat` command indicates the primary path is no longer available. The secondary path (vmhba3) is still available and now an active path, as indicated by the asterisk (*). The primary path vmhba2 is the preferred primary path, indicated by the pound sign (#), even though it is temporarily unavailable.

Note: An active virtual device and the asterisk (*) (i.e., vmhba3:0:1 on*) may or may not mean there is I/O running. The asterisk only indicates a complete failover or failback state of a device.

A complete display of a failover path can be viewed from the Options of Storage Management of VMware Management Interface (see Figure 5.2).

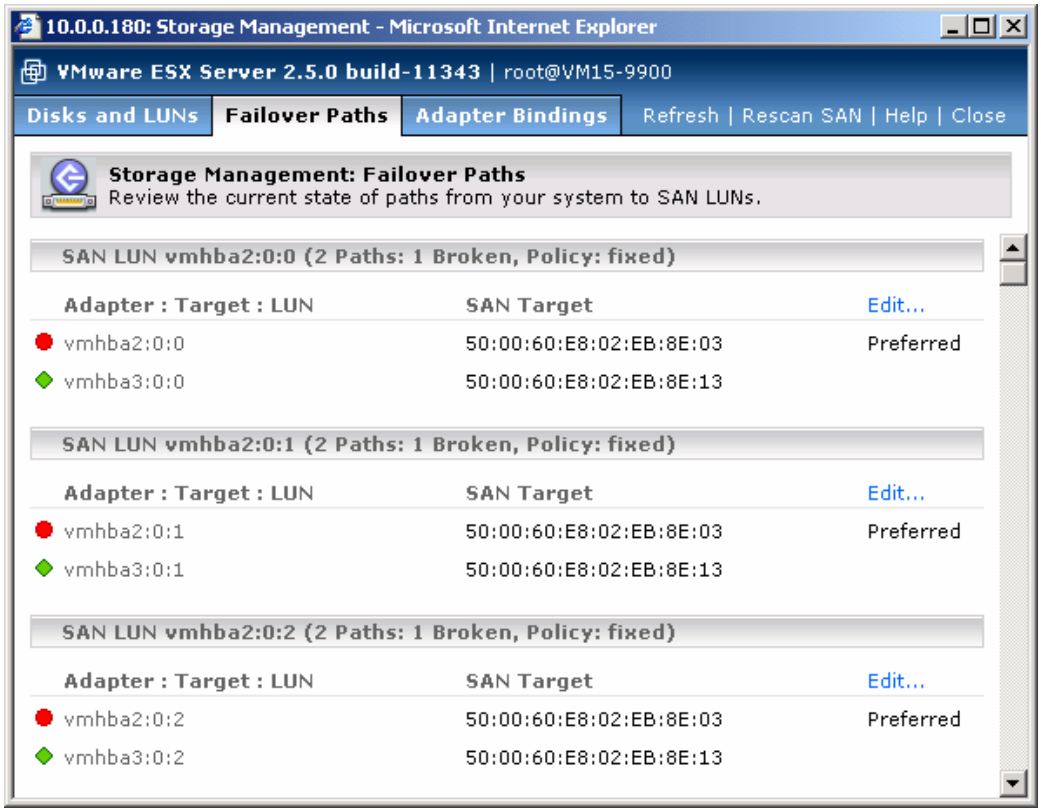

Figure 5.2 Failover Paths of VMware Management Interface

VMware with the preferred and fixed policy supports auto-failback. Presumably, a primary device path was in a good active state with a preferred fixed policy has failed. It automatically becomes a preferred path when the physical path is re-connected. If there is I/O, the preferred path becomes active. You can use the `ls` and `cat` commands to verify the complete failback state of both primary and secondary paths (see Figure 5.3).

```
[… root]# cat /proc/VMware/scsi/vmhba2/0:0

Paths:fixed
      vmhba2:0:0 on#
      vmhba3:0:0 on*


[… root]# cat /proc/VMware/scsi/vmhba2/0:1

Paths:fixed
      vmhba2:0:1 on#
      vmhba3:0:1 on*
```

Figure 5.3 Path Failback Display

You can use the `vmkmultipath –q` command to display multiple entries of failover devices of all vmhbas. The information includes device capacity and explicit failover state for each virtual device. Figure 5.4 shows the output of the command:

```
[… root]# vmkmultipath –q | more

Disk and multipath information follows

Disk vmhba2:0:0 (34,757 MB) has 2 paths. Policy is fixed.
     vmhba2:0:0      on      (active, preferred)
        vmhba3:0:0   on

Disk vmhba2:0:1 (34,757 MB) has 2 paths. Policy is fixed.
     vmhba2:0:1      on      (active, preferred)
        vmhba3:0:1   on

Disk vmhba2:0:2 (34,757 MB) has 2 paths. Policy is fixed.
     vmhba2:0:2      on      (active, preferred)
        vmhba3:0:2   on

Disk vmhba2:0:3 (34,757 MB) has 2 paths. Policy is fixed.
     vmhba2:0:3      on      (active, preferred)
        vmhba3:0:3   on

Disk vmhba2:0:4 (34,757 MB) has 2 paths. Policy is fixed.
     vmhba2:0:4      on      (active, preferred)
        vmhba3:0:4   on


.
.
.


Disk vmhb4:0:0 (4,693 MB) has 2 paths. Policy is fixed.
     Vmhba4:0:0      on      (active, preferred)
        Vmhba5:0:0   on

Disk vmhb4:0:1 (4,693 MB) has 2 paths. Policy is fixed.
     Vmhba4:0:1      on      (active, preferred)
        Vmhba5:0:1   on

Disk vmhb4:0:2 (4,693 MB) has 2 paths. Policy is fixed.
     Vmhba4:0:2      on      (active, preferred)
        Vmhba5:0:2   on

Disk vmhb4:0:3 (4,693 MB) has 2 paths. Policy is fixed.
     Vmhba4:0:3      on      (active, preferred)
        Vmhba5:0:3   on

Disk vmhb4:0:4 (4,693 MB) has 2 paths. Policy is fixed.
     Vmhba4:0:4      on      (active, preferred)
        Vmhba5:0:4   on

.
.
.
```

Figure 5.4 Vmkmultipath –q Command

5.3 Fixed Policy with Preferred Paths v.s. MRU

By default, VMware sets all virtual devices with a fixed policy. The vmhba with the lower bus ID number (i.e., vmhba2 v.s. vmhba3) is chosen as the preferred (primary) path. The other path that shares the same virtual device is the secondary path. The order number of the vmhba depends on the result of the VMware scanning of the newly added HBA on the PCI slot (slot1, slot2, etc.), with the existing vmhba on the ESX Server PCI slot. For complete information about how VMware determines which vmhba has a lower ID number, refer to the ESX Server Administration Guide on VMware website http://www.VMware.com/support/pubs/esx_pubs.html

- **Fixed with Preferred Path:** Fixed policy with Preferred Path is an active/passive failover mechanism of VMware ESX Server. One path is chosen as a preferred (primary) path. The secondary path becomes active if the primary path is no longer available. The preferred path is the primary path and always preferred in the VMware ESX Server, so long as the path is viable and there is I/O.
- **MRU:** MRU is the other active/passive failover mechanism of VMware ESX Server 2.5. There is no pre-selected primary path for the MRU. All shared paths to the virtual device are primary paths. Any shared path in an MRU policy can be determined as an active path and always an active path as long as it is healthy. An MRU policy behaves like a fixed policy without auto-failback feature. However, a connected path can return to an active path if it is manually set as an active one again, after failure.

Note: Fixed with Preferred Path and MRU policy can be set at the device level, not at the vmhba level.

Execute the following command at the VMware service console to change a Preferred Fixed policy to a MRU policy for a virtual device.

```
[… root]# vmkmultipath –s vmhba2:0:0 –p mru

VMware::ExtHelpers::System[788] '/sbin/fdisk' –l /dev/sda


[… root]# vmkmultipath –q vmhba2:0:0
Disk and multipath information follows:

Disk vmhba2:0:0 (34,757 MB) has 2 paths. Policy is mru.
     vmhba2:0:0      on      (active, preferred)
     vmhba3:0:0      on
```

Figure 5.5 Changing Fixed Policy

Figure 5.6 and Figure 5.7 show how to change a fixed policy to MRU policy and select a preferred path within a fixed policy using the **Edit** field of the **Failover Paths** menu.

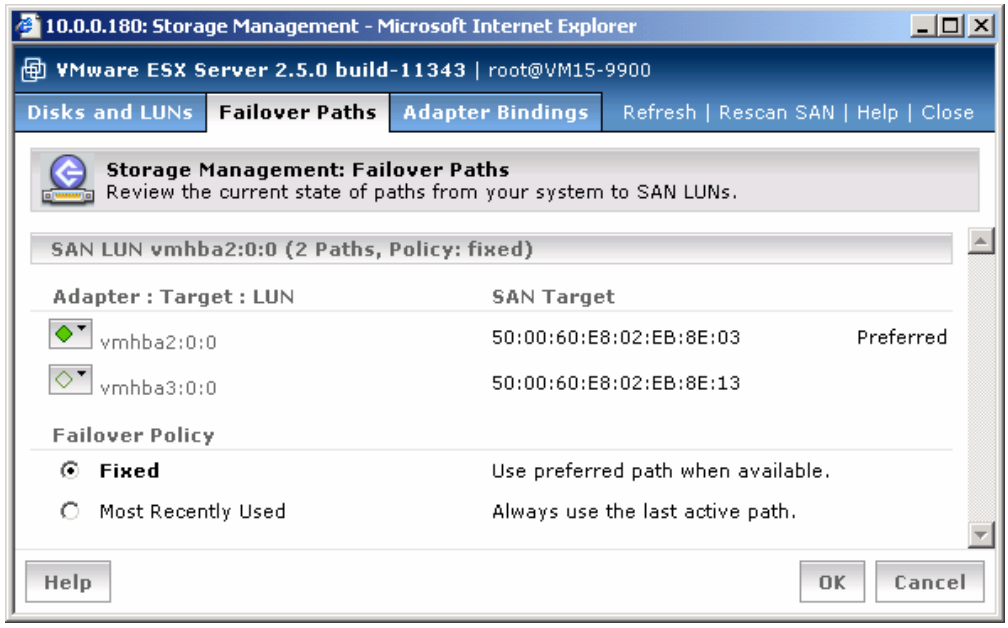

Figure 5.6 Changing Fixed Policy to MRU

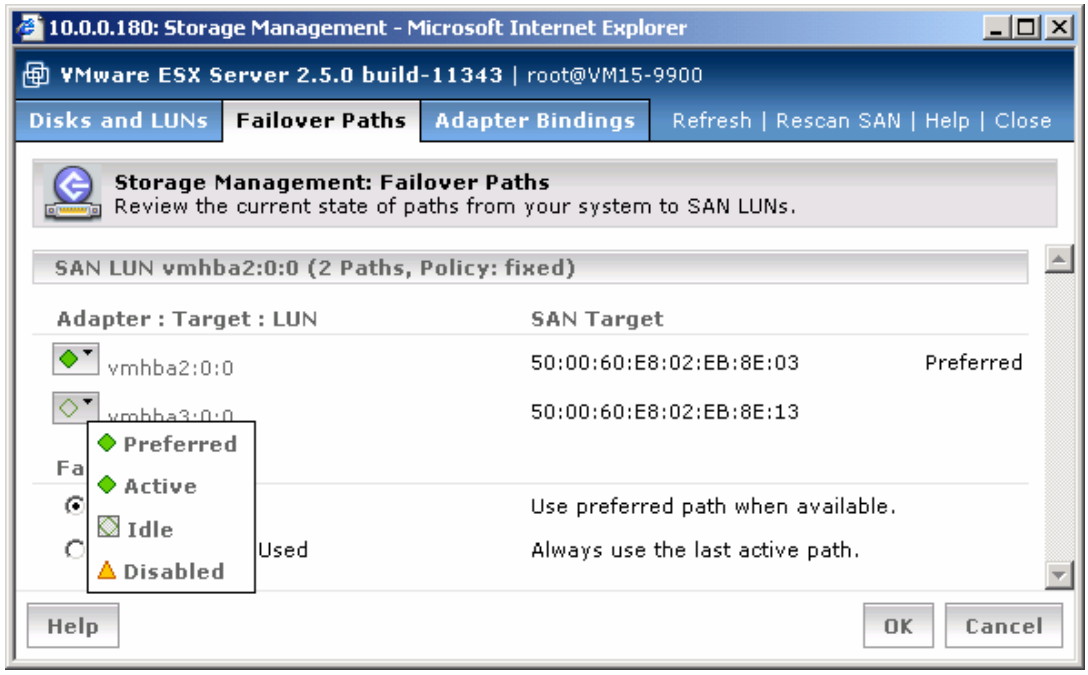

Figure 5.7 Selecting Preferred Path of Fixed Policy

Chapter 6 Troubleshooting

The AMS/WMS provides continuous data availability. For troubleshooting information for the AMS/WMS system, refer to the following documents:

- *Hitachi TagmaStore Adaptable Modular Storage 200 User and Reference Guide*, MK-95DF713
- *Hitachi TagmaStore Adaptable Modular Storage 500 User and Reference Guide*, MK-95DF714
- *Hitachi TagmaStore Adaptable Modular Storage 1000 User and Reference Guide*, MK-95DF780
- *Hitachi TagmaStore Workgroup Modular Storage 100 User and Reference Guide, MK-95DF738*

Table 6.1 lists potential error conditions during AMS/WMS VMware device configuration and provides instructions for resolving each condition. If you are unable to resolve an error condition, please consult your Hitachi Data Systems representative, or call the Hitachi Data Systems Support Center for assistance. See section 6.1 for instructions on calling the Hitachi Data Systems Support Center.

Table 6.1 Troubleshooting (continues on following page)

Error Condition	Recommended Action
FC HBA does not show log-in LED indicators	Make sure the FC cable is connected to a operational FC switch Other possible cause includes FC HBA drivers are not loaded. Type vmkload_mod –l to verify QLogic qla2300_607.o and/or Emulex lpfcdd_2xx.o has loaded. There is a conflict of mixing of different type of HBA in the system
FC HBA driver is not loaded	Possible cause is that the vmhba resource has not been properly assigned as dedicated to a virtual machine in the Startup Profile. Properly assigned HBA to a virtual machine in the Startup Profile, followed by an ESX Server reboot will automatically load the HBA driver by VMware
VMware Management Interface fails to access VMware ESX Server log-on screen	Make sure there is IP communication between the VMware Management Interface and the ESX Server. Httpd service may not have been started. Verify by typing "ps –ef |more" at the service console to check for httpd services. If not started, gracefully reboot ESX Server by typing "shutdown –h now"
511 Error	Most likely there is a SCSI device error timing-out. Check for such error messages in the log file. Another possible cause could just be a timing-out of a GUI. Restart the GUI and restart the operation.
Cannot format an NTFS or mkfs a device successfully	Possibly an intermittent cabling problem or a bad storage port.
Virtual Machine HBA does not see Lun8 and greater	Verify cabling, storage LUN, switch and storage security and LUN masking. Check Disk.MaxLUN parameter in the Advance Settings (VMware Management Interface) is set to greater than 8.

Table 6.1 Troubleshooting (continued)

Error Condition	Recommended Action
Vmkfstools –s does not add LUN online	Delete the LUN. Select and add another LUN and retry the process again. Repeat the command and/or perform the Rescan SAN function in the Storage Management of the VMware Management Interface and display again
Service console discover online LUN addition but the Disks and LUNs does not	Rescan SAN and refresh
VMware ESX Server crashes as booting up	Check for the error message on the screen. It could be because of mixing different types of HBA in the server.
VMware does not show * for failover or failback device	It is possible that there is no I/O to the device. A rescan and refresh with the Failover Paths of the Storage Management may update the display status.
Creating VMFS-2 file shows vmfs-2.11	ESX Server 2.5 still uses VMFS-2.11
Guest OS virtual machine booting up but not installing the OS	It is possible that there is an existing corrupted vmdk file (usually because of previous incomplete installation). Delete the vmdk file from the File Manager and remove it from the Guest OS. Add a new device for the Guest OS and recreate a new vmdk image file.
Cannot add Meta Data File for raw device	The Meta Data File for the raw device may have existed. Selected the existing Meta Data File or delete the old Meta Data File and create a new one.
Cannot create vmdk mapping file for raw device	Make sure vmkfs file system is vmkfs2.11 and the raw device attached to a non-shared vmhba (dedicated)
Volume label is not successful	Limit the number of characters to 30
mkfs –t ext3 does not create ext3 file system for Linux 2.1	Use mkfs –j for Linux 2.1
Cannot delete a VMFS file	It is possible that there is an active swap file on the same extended partition. Manually turn off the swap device (using `vmkfstools` command) from the service console and try again. Relocate the swap file to another disk.
Guest OS cannot communicate with the server or outside network	Make sure a virtual switch is created and bound to a connected network adapter

6.1 Calling the Support Center

If you need to call Hitachi Data Systems Support Center, make sure that you provide as much information about the problem as possible, including circumstances surrounding error or failure and the exact content of error messages displayed on the host system(s).

The worldwide Hitachi Data Systems Support Centers are:

- Hitachi Data Systems North America/Latin America
  San Diego, California, USA
  1-800-446-0744
- Hitachi Data Systems Europe
  Contact Hitachi Data Systems Local Support
- Hitachi Data Systems Asia Pacific
  North Ryde, Australia
  61-2-9325-3300

Appendix A VMware Command Usage

For VMware 2.5, a few VMware commands such as `vmkfstools` and `vmkmultipath` are used for various configuration and display activities. For the complete description of how to use these commands, refer to the VMware ESX Server Administration Guide.

Acronyms and Abbreviations

AL	arbitrated loop
AL-PA	arbitrated loop physical address
AMS	TagmaStore Adaptable Modular Storage
ESCON®	Enterprise System Connection (IBM registered trademark for optical channels)
FC	fibre-channel
FC-AL	fibre-channel arbitrated loop
GB	Gigabytes
GUI	graphical user interface
HBA	host bus adapter
I/O	input/output
LDEV	logical device
LU	logical unit
LUN	logical unit, logical unit number
LVI	Logical Volume Image
MB	Megabytes
MRU	most recently used
OFC	open fibre control
OS	operating system
PCI	power control interface
RAID	redundant array of independent disks
SAN	Storage Area Network
SCSI	small computer system interface
TID	target ID
vmhba	virtual machine host bus adapter
WMS	TagmaStore Workgroup Modular Storage