ESX Server 3 Configuration Guide

Update 2 and later for
ESX Server 3.5 and VirtualCenter 2.5

This document supports the version of each product listed and supports all subsequent versions until the document is replaced by a new edition. To check for more recent editions of this document, see http://www.vmware.com/support/pubs.

EN-000031-04

vmware®

You can find the most up-to-date technical documentation on the VMware Web site at:

http://www.vmware.com/support/

The VMware Web site also provides the latest product updates.

If you have comments about this documentation, submit your feedback to:

docfeedback@vmware.com

VMware, Inc.
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Contents

About This Book 9

1 Introduction 13
- Networking 14
- Storage 14
- Security 15
- Appendixes 15

Networking

2 Networking 19
- Networking Concepts Overview 20
  - Virtual Switches 21
  - Port Groups 24
- Enabling Network Services 24
- Viewing Networking Information in the VI Client 25
- Virtual Network Configuration for Virtual Machines 27
- VMkernel Networking Configuration 30
  - TCP/IP Stack at the VMkernel Level 30
  - Implications and Guidelines for Configuration 31
- Service Console Configuration 34
  - Basic Service Console Configuration Tasks 34
  - Using DHCP for the Service Console 39

3 Advanced Networking 41
- Virtual Switch Properties and Policies 42
  - Virtual Switch Properties 42
  - Virtual Switch Policies 50
- Port Group Configuration 58
- DNS and Routing 61

TCP Segmentation Offload and Jumbo Frames 62
Enabling TSO 62
Enabling Jumbo Frames 64
NetQueue and Networking Performance 65
Setting Up MAC Addresses 67
MAC Addresses Generation 67
Setting MAC Addresses 68
Using MAC Addresses 69
Networking Best Practices and Tips 69
Networking Best Practices 69
Networking Tips 70

4 Networking Scenarios and Troubleshooting 73
Networking Configuration for Software iSCSI Storage 73
Configuring Networking on Blade Servers 78
Troubleshooting 81
Troubleshooting Service Console Networking 81
Troubleshooting Network Adapter Configuration 82
Troubleshooting Physical Switch Configuration 82
Troubleshooting Port Group Configuration 83

Storage

5 Introduction to Storage 87
Storage Overview 88
Types of Physical Storage 88
Local Storage 89
Networked Storage 90
Supported Storage Adapters 91
Datastores 91
VMFS Datastores 92
NFS Datastore 95
How Virtual Machines Access Storage 95
Comparing Types of Storage 97
Viewing Storage Information in the VMware Infrastructure Client 97
Displaying Datastores 98
Viewing Storage Adapters 99
Understanding Storage Device Naming in the Display 100
Configuring and Managing Storage 101

6 Configuring Storage 103
Local Storage 104
Adding Local Storage 104
Fibre Channel Storage 107
Adding Fibre Channel Storage 108
iSCSI Storage 110
iSCSI Initiators 110
Naming Requirements 111
Discovery Methods 112
iSCSI Security 112
Configuring Hardware iSCSI Initiators and Storage 113
Configuring Software iSCSI Initiators and Storage 121
Performing a Rescan 126
Network Attached Storage 127
How Virtual Machines Use NFS 127
NFS Volumes and Virtual Machine Delegate Users 129
Configuring ESX Server 3 to Access NFS Volumes 129
Creating an NFS-Based Datastore 129
Creating a Diagnostic Partition 130

7 Managing Storage 133
Managing Datastores 133
Editing VMFS Datastores 134
Upgrading Datastores 135
Changing the Names of Datastores 136
Adding Extents to Datastores 136
Managing Multiple Paths 137
Multipathing with Local Storage and Fibre Channel SANs 138
Multipathing with iSCSI SAN 140
Viewing the Current Multipathing Status 141
Setting Multipathing Policies for LUNs 143
Disabling Paths 144
The vmkfstools Commands 144

8 Raw Device Mapping 145
About Raw Device Mapping 145
Benefits of Raw Device Mapping 147
Limitations of Raw Device Mapping 149

Raw Device Mapping Characteristics 150
Virtual Compatibility Mode Compared to Physical Compatibility Mode 151
Dynamic Name Resolution 152
Raw Device Mapping with Virtual Machine Clusters 154
Comparing Raw Device Mapping to Other Means of SCSI Device Access 154
Managing Mapped LUNs 155
VMware Infrastructure Client 155
The vmkfstools Utility 158
File System Operations 159

Security

9 Security for ESX Server 3 Systems 163
ESX Server 3 Architecture and Security Features 164
Security and the Virtualization Layer 164
Security and Virtual Machines 165
Security and the Service Console 167
Security and the Virtual Networking Layer 169
Security Resources and Information 175

10 Securing an ESX Server 3 Configuration 177
Securing the Network with Firewalls 177
Firewalls for Configurations with a VirtualCenter Server 179
Firewalls for Configurations Without a VirtualCenter Server 182
TCP and UDP Ports for Management Access 183
Connecting to VirtualCenter Server Through a Firewall 185
Connecting to the Virtual Machine Console Through a Firewall 186
Connecting ESX Server 3 Hosts Through Firewalls 188
Opening Firewall Ports for Supported Services and Management Agents 188
Securing Virtual Machines with VLANs 195
Security Considerations for vSwitches and VLANs 198
Virtual Switch Protection and VLANs 200
Securing Virtual Switch Ports 202
Securing iSCSI Storage 204
Securing iSCSI Devices Through Authentication 205
Protecting an iSCSI SAN 208

11 Authentication and User Management 211
Securing ESX Server 3 Through Authentication and Permissions 211
About Users, Groups, Permissions, and Roles 213
Working with Users and Groups on ESX Server 3 Hosts 219
Encryption and Security Certificates for ESX Server 3 225
Adding Certificates and Modifying ESX Server 3 Web Proxy Settings 227
Regenerating Certificates 231
Replacing Self-Signed Certificates with CA-Signed Certificates 231
Virtual Machine Delegates for NFS Storage 232

12 Service Console Security 235
General Security Recommendations 236
Logging On to the Service Console 237
Service Console Firewall Configuration 237
Changing the Service Console Security Level 238
Opening and Closing Ports in the Service Console Firewall 240
Password Restrictions 241
Password Aging 242
Password Complexity 244
Changing the Password Plug-In 248
Cipher Strength 249
setuid and setgid Applications 250
Default setuid Applications 251
Default setgid Applications 252
SSH Security 253
Security Patches and Security Vulnerability Scanning Software 254

13 Security Deployments and Recommendations 257
Security Approaches for Common ESX Server 3 Deployments 257
Single Customer Deployment 257
Multiple Customer Restricted Deployment 259
Multiple Customer Open Deployment 261
Virtual Machine Recommendations 263
Installing Antivirus Software 263
Disabling Copy and Paste Operations Between the Guest Operating System and Remote Console 263
Removing Unnecessary Hardware Devices 265
Limiting Guest Operating System Writes to Host Memory 267
Configuring Logging Levels for the Guest Operating System 270

Appendixes

A ESX Server 3 Technical Support Commands 277
Other Commands 282

B Using vmkfstools 283
vmkfstools Command Syntax 283
vmkfstools Options 284
-v Suboption 285
File System Options 285
Virtual Disk Options 288
Managing SCSI Reservations of LUNs 294

Index 297

About This Book

This book, the *ESX Server 3 Configuration Guide*, provides information on how to configure networking for ESX Server 3, including how to create virtual switches and ports and how to set up networking for virtual machines, VMotion, IP storage, and the service console. It also discusses configuring file system and various types of storage such as iSCSI, Fibre Channel, and so forth. To help you protect your ESX Server 3 installation, the guide provides a discussion of security features built into ESX Server 3 and the measures you can take to safeguard it from attack. In addition, it includes a list of ESX Server 3 technical support commands along with their VI Client equivalents and a description of the `vmkfstools` utility.

The *ESX Server 3 Configuration Guide* covers ESX Server 3.5. To read about ESX Server 3i version 3.5, see http://www.vmware.com/support/pubs/vi_pubs.html.

For ease of discussion, this book uses the following product naming conventions:

- For topics specific to ESX Server 3.5, this book uses the term "ESX Server 3."
- For topics specific to ESX Server 3i version 3.5, this book uses the term "ESX Server 3i."
- For topics common to both products, this book uses the term "ESX Server."
- When the identification of a specific release is important to a discussion, this book refers to the product by its full, versioned name.

When a discussion applies to all versions of ESX Server for VMware® Infrastructure 3, this book uses the term "ESX Server 3.x."

Intended Audience

This book is intended for anyone who wants to install, upgrade, or use ESX Server 3. The information in this book is written for experienced Windows or Linux system administrators who are familiar with virtual machine technology and datacenter operations.

Document Feedback

VMware welcomes your suggestions for improving our documentation. If you have comments, send your feedback to:

docfeedback@vmware.com

VMware Infrastructure Documentation

The VMware Infrastructure documentation consists of the combined VMware VirtualCenter and ESX Server documentation set.

Abbreviations Used in Figures

The figures in this book use the abbreviations listed in Table 1.

Table 1. Abbreviations

Abbreviation	Description
database	VirtualCenter database
datastore	Storage for the managed host
dsk#	Storage disk for the managed host
host*n*	VirtualCenter managed hosts
SAN	Storage area network type datastore shared between managed hosts
tmplt	Template
user#	User with access permissions
VC	VirtualCenter
VM#	Virtual machines on a managed host
VM	Virtual machine
VI Client	VMware Infrastructure Client
server	VirtualCenter Server

Technical Support and Education Resources

The following sections describe the technical support resources available to you. To access the current versions of this book and other books, go to:

http://www.vmware.com/support/pubs.

Online and Telephone Support

Use online support to submit technical support requests, view your product and contract information, and register your products. Go to:

http://www.vmware.com/support

Customers with appropriate support contracts should use telephone support for the fastest response on priority 1 issues. Go to:

http://www.vmware.com/support/phone_support.html

Support Offerings

Find out how VMware support offerings can help meet your business needs. Go to:

http://www.vmware.com/support/services

VMware Professional Services

VMware Education Services courses offer extensive hands-on labs, case study examples, and course materials designed to be used as on-the-job reference tools. Courses are available onsite, in the classroom, and live online. For onsite pilot programs and implementation best practices, VMware Consulting Services provides offerings to help you assess, plan, build, and manage your virtual environment. To access information about education classes, certification programs, and consulting services, go to http://www.vmware.com/services.

1 Introduction

The *ESX Server 3 Configuration Guide* describes the tasks you need to complete to configure ESX Server 3 host networking, storage, and security. In addition, it provides overviews, recommendations, and conceptual discussions to help you understand these tasks and how to deploy an ESX Server 3 host to meet your needs. Before you use the information in the *ESX Server 3 Configuration Guide*, read the *Introduction to Virtual Infrastructure* for an overview of system architecture and the physical and virtual devices that make up a VMware Infrastructure system.

This introduction summarizes the contents of this guide so that you can find the information you need. This guide discusses these subjects:

- ESX Server 3 network configurations
- ESX Server 3 storage configurations
- ESX Server 3 security features
- ESX Server 3 command reference
- The `vmkfstools` command

Networking

The ESX Server 3 networking chapters provide you with a conceptual understanding of physical and virtual network concepts, a description of the basic tasks you need to complete to configure your ESX Server 3 host's network connections, and a discussion of advanced networking topics and tasks. The networking section contains the following chapters:

- Networking – Introduces you to network concepts and guides you through the most common tasks you need to complete when setting up the network for the ESX Server 3 host.
- Advanced Networking – Discusses advanced networking tasks such as setting up MAC addresses, editing virtual switches and ports, and DNS routing. In addition, it provides tips on making your network configuration more efficient.
- Networking Scenarios and Troubleshooting – Describes common networking configuration and troubleshooting scenarios.

Storage

The ESX Server 3 storage chapters provide you with a basic understanding of storage, a description of the basic tasks you perform to configure and manage your ESX Server 3 host's storage, and a discussion of how to set up raw device mapping (RDM). The storage section contains the following chapters:

- Introduction to Storage – Introduces you to the types of storage you can configure for the ESX Server 3 host.
- Configuring Storage – Explains how to configure local SCSI storage, Fibre Channel storage, and iSCSI storage. It also addresses virtual machine file system (VMFS) storage and network-attached storage.
- Managing Storage – Explains how to manage existing datastores and the file systems that comprise datastores.
- Raw Device Mapping – Discusses raw device mapping, how to configure this type of storage, and how to manage raw device mappings by setting up multipathing, failover, and so forth.

Security

The ESX Server 3 security chapters discuss safeguards that VMware has built into ESX Server 3 and measures you can take to protect your ESX Server 3 host from security threats. These measures include using firewalls, taking advantage of the security features of virtual switches, and setting up user authentication and permissions. The security section contains the following chapters:

- Security for ESX Server 3 Systems – Introduces you to the ESX Server 3 features that help you ensure a secure environment for your data and gives you an overview of system design as it relates to security.
- Securing an ESX Server 3 Configuration – Explains how to configure firewall ports for ESX Server 3 hosts and VMware VirtualCenter, how to use virtual switches and VLANs to ensure network isolation for virtual machines, and how to secure iSCSI storage.
- Authentication and User Management – Discusses how to set up users, groups, permissions, and roles to control access to ESX Server 3 hosts and VirtualCenter. It also discusses encryption and delegate users.
- Service Console Security – Discusses the security features built into the service console and shows you how to configure these features.
- Security Deployments and Recommendations – Provides some sample deployments to give you an idea of the issues you need to consider when you set up your own ESX Server 3 deployment. This chapter also tells you about actions you can take to further secure virtual machines.

Appendixes

The *ESX Server 3 Configuration Guide* includes appendixes that provide specialized information you might find useful when configuring an ESX Server 3 host.

- ESX Server 3 Technical Support Commands – Discusses the ESX Server 3 configuration commands that you can issue through a command-line shell such as secure shell (SSH). Although these commands are available for your use, do not consider them to be an API that you can build scripts on. These commands are subject to change and VMware does not support applications and scripts that rely on ESX Server 3 configuration commands. This appendix provides you with VMware Infrastructure Client equivalents for these commands.
- Using vmkfstools – Discusses the `vmkfstools` utility, which you can use to perform management and migration tasks for iSCSI disks.

Networking

Networking 2

This chapter guides you through the basic concepts of networking in the ESX Server 3 environment and how to set up and configure a network in a virtual infrastructure environment.

Use the VMware Infrastructure (VI) Client to add networking based on three categories that reflect the three types of network services:

- Virtual machines
- VMkernel
- Service console

This chapter discusses the following topics:

- "Networking Concepts Overview" on page 20
- "Enabling Network Services" on page 24
- "Viewing Networking Information in the VI Client" on page 25
- "Virtual Network Configuration for Virtual Machines" on page 27
- "VMkernel Networking Configuration" on page 30
- "Service Console Configuration" on page 34

Networking Concepts Overview

A few concepts are essential to a thorough understanding of virtual networking. If you are new to ESX Server 3, VMware recommends that you read this section.

A *physical network* is a network of physical machines that are connected so that they can send data to and receive data from each other. VMware ESX Server 3 runs on a physical machine.

A *virtual network* is a network of virtual machines running on a single physical machine that are connected logically to each other so that they can send data to and receive data from each other. Virtual machines can be connected to the virtual networks that you create in the procedure to add a network. Each virtual network is serviced by a single virtual switch. A virtual network can be connected to a physical network by associating one or more physical Ethernet adapters, also referred to as uplink adapters, with the virtual network's virtual switch. If no uplink adapters are associated with the virtual switch, all traffic on the virtual network is confined within the physical host machine. If one or more uplink adapters are associated with the virtual switch, virtual machines connected to that virtual network can also access the physical networks connected to the uplink adapters.

A *physical Ethernet switch* manages network traffic between machines on the physical network. A switch has multiple ports, each of which can be connected to a single other machine or another switch on the network. Each port can be configured to behave in certain ways depending on the needs of the machine connected to it. The switch learns which hosts are connected to which of its ports and uses that information to forward traffic to the correct physical machines. Switches are the core of a physical network. Multiple switches can be connected together to form larger networks.

A virtual switch, *vSwitch*, works much like a physical Ethernet switch. It detects which virtual machines are logically connected to each of its virtual ports and uses that information to forward traffic to the correct virtual machines. A vSwitch can be connected to physical switches by using physical Ethernet adapters, also referred to as uplink adapters, to join virtual networks with physical networks. This type of connection is similar to connecting physical switches together to create a larger network. Even though a vSwitch works much like a physical switch, it does not have some of the advanced functionality of a physical switch. See "Virtual Switches" on page 21.

A *port group* specifies port configuration options such as bandwidth limitations and VLAN tagging policies for each member port. Network services connect to vSwitches through port groups. Port groups define how a connection is made through the vSwitch to the network. In typical use, one or more port groups is associated with a single vSwitch. See "Port Groups" on page 24.

NIC teaming occurs when multiple uplink adapters are associated with a single vSwitch to form a team. A team can either share the load of traffic between physical and virtual networks among some or all of its members, or provide passive failover in the event of a hardware failure or a network outage.

VLANs enable a single physical LAN segment to be further segmented so that groups of ports are isolated from one another as if they were on physically different segments. The standard is 802.1Q.

The *VMkernel TCP/IP networking stack* supports iSCSI, NFS, and VMotion. Virtual machines run their own systems' TCP/IP stacks, and connect to the VMkernel at the Ethernet level through virtual switches. Two new features in ESX Server 3, iSCSI and NFS, are referred to as *IP storage* in this chapter. IP storage refers to any form of storage that uses TCP/IP network communication as its foundation. iSCSI can be used as a virtual machine datastore, and NFS can be used as a virtual machine datastore and for direct mounting of `.ISO` files, which are presented as CD-ROMs to virtual machines.

NOTE The networking chapters discuss how to set up networking for iSCSI and NFS. To configure the storage portion of iSCSI and NFS, see the storage chapters.

TCP Segmentation Offload, *TSO*, allows a TCP/IP stack to emit very large frames (up to 64KB) even though the maximum transmission unit (MTU) of the interface is smaller. The network adapter then separates the large frame into MTU-sized frames and prepends an adjusted copy of the initial TCP/IP headers. See "TCP Segmentation Offload and Jumbo Frames" on page 62.

Migration with VMotion enables a virtual machine that is powered on to be transferred from one ESX Server 3 host to another without shutting down the virtual machine. The optional VMotion feature requires its own license key.

Virtual Switches

VMware Infrastructure lets you use the Virtual Infrastructure (VI) Client or direct SDK APIs to create abstracted network devices called virtual switches (vSwitches). A vSwitch can route traffic internally between virtual machines and link to external networks.

NOTE You can create a maximum of 127 vSwitches on a single host.

Use virtual switches to combine the bandwidth of multiple network adapters and balance communications traffic among them. They can also be configured to handle physical NIC failover.

A vSwitch models a physical Ethernet switch. The default number of logical ports for a vSwitch is 56. However, a vSwitch can be created with up to 1016 ports in ESX Server 3. You can connect one network adapter of a virtual machine to each port. Each uplink adapter associated with a vSwitch uses one port. Each logical port on the vSwitch is a member of a single port group. Each vSwitch can also have one or more port groups assigned to it. See "Port Groups" on page 24.

Before you can configure virtual machines to access a network, you must perform the following tasks:

1 Create a vSwitch, and configure it to connect to the physical adapters on the host for the required physical network.

2 Create a virtual machine port group connected to that vSwitch, and give it a name that will be referenced by the virtual machine configuration.

When two or more virtual machines are connected to the same vSwitch, network traffic between them is routed locally. If an uplink adapter is attached to the vSwitch, each virtual machine can access the external network that the adapter is connected to, as shown in Figure 2-1.

Figure 2-1. Virtual Switch Connections

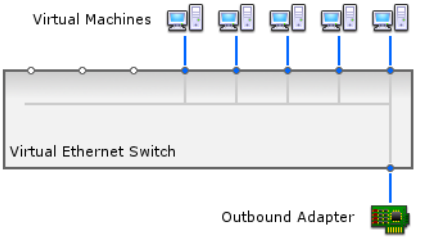

In the VI Client, the details for the selected vSwitch are presented as an interactive diagram as shown in Figure 2-2. The most important information for each vSwitch is always visible.

Figure 2-2. Virtual Switch Interactive Diagram

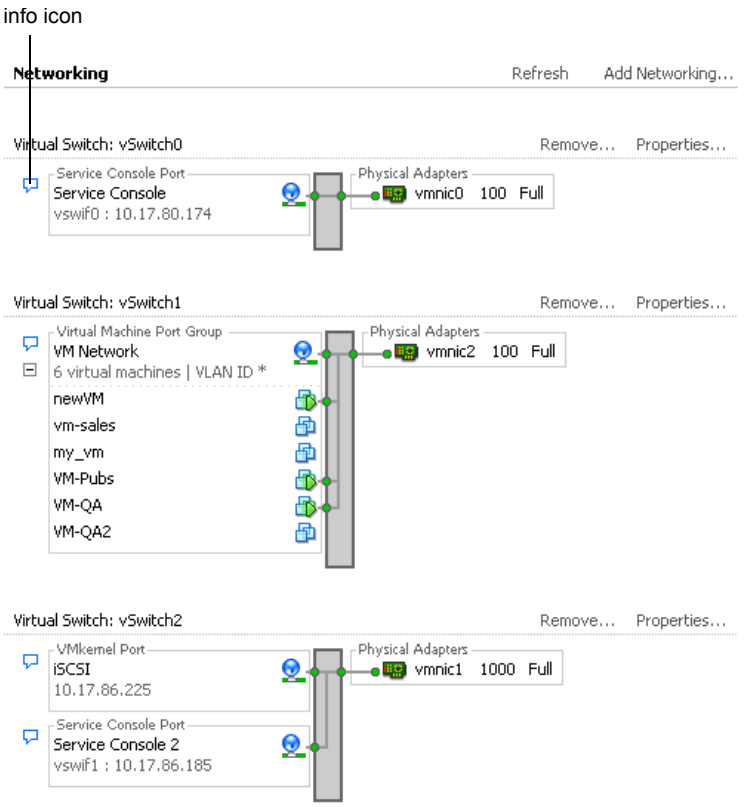

Click the info icon to selectively reveal secondary and tertiary information.

A pop-up window displays detailed properties, as shown in Figure 2-3.

Figure 2-3. Virtual Switch Detailed Properties

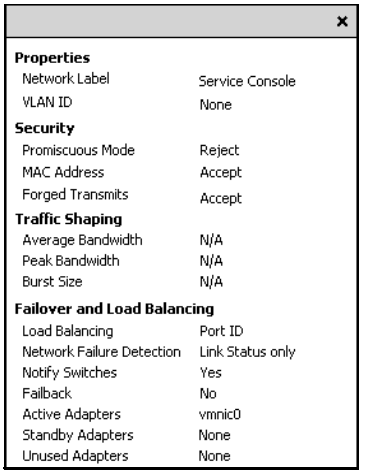

Port Groups

Port groups aggregate multiple ports under a common configuration and provide a stable anchor point for virtual machines connecting to labeled networks. Each port group is identified by a network label, which is unique to the current host.

NOTE You can create a maximum of 512 port groups on a single host.

A VLAN ID, which restricts port group traffic to a logical Ethernet segment within the physical network, is optional.

NOTE For a port group to reach port groups located on other VLANs, set the VLAN ID to 4095.

Enabling Network Services

You need to enable two types of network services in ESX Server 3:

- Connecting virtual machines to the physical network
- Connecting VMkernel services (such as NFS, iSCSI, or VMotion) to the physical network

Networking for the service console, which runs management services for ESX Server 3, is set up by default during installation. A service console port is required for ESX Server 3 to connect to any network or remote services, including the VI Client. Additional service console ports might be necessary for certain services, such as iSCSI storage. For information on configuring service console ports, see "Service Console Configuration" on page 34.

Viewing Networking Information in the VI Client

The VI Client displays general networking information and information specific to network adapters.

To view general networking information in the VI Client

1 Log in to the VI Client and select the server from the inventory panel.

The hardware configuration page for this server appears.

2 Click the **Configuration** tab and click **Networking**.

The networking panel displays the following information, as shown in Figure 2-4:

- Virtual switches
- Adapter information for each adapter
  - Link status
  - Apparent speed and duplex
- Service console and VMkernel TCP/IP services

  IP address
- Service console

  Virtual device name
- Virtual machines
  - Power status
  - Connection status
- Port group
  - Network label – common to all three port configuration types
  - Number of configured virtual machines
  - VLAN ID, if any – common to all three port configuration types

Figure 2-4. General Networking Information

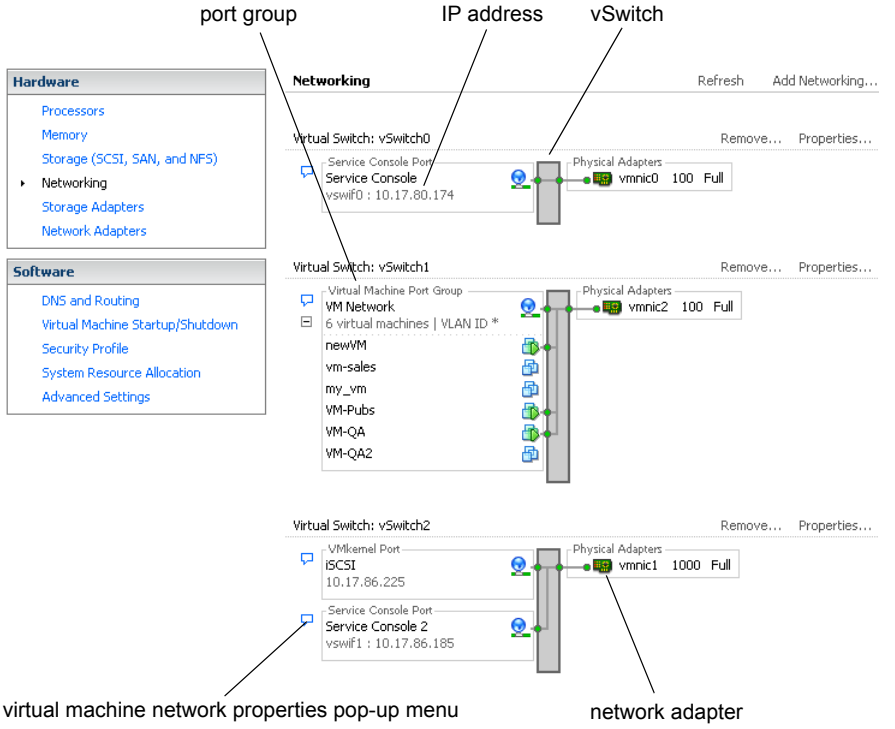

To view network adapter information in the VI Client

1 Log in to the VI Client and select the server from the inventory panel.

The hardware configuration page for this server appears.

2 Click the **Configuration** tab, and click **Network Adapters**.

The network adapters panel displays the following information:

- **Device** – Name of the network adapter
- **Speed** – Actual speed and duplex of the network adapter
- **Configured** – Configured speed and duplex of the network adapter
- **vSwitch** – vSwitch that the network adapter is associated with
- **Observed IP ranges** – IP addresses that the network adapter has access to
- **Wake on LAN supported** – Network adapter ability to support Wake on LAN

Virtual Network Configuration for Virtual Machines

The VI Client Add Network wizard guides you through the tasks to create a virtual network to which virtual machines can connect. These tasks include:

- Setting the connection type for a virtual machine
- Adding the virtual network to a new or an existing vSwitch
- Configuring the connection settings for the network label and the VLAN ID

For information on configuring network connections for an individual virtual machine, see the *Basic System Administration Guide*.

When you set up virtual machine networks, consider whether you want to migrate the virtual machines in the network between ESX Server 3 hosts. If so, be sure that both hosts are in the same broadcast domain—that is, the same Layer 2 subnet.

ESX Server 3 doesn't support virtual machine migration between hosts in different broadcast domains because the migrated virtual machine might require systems and resources that it would no longer have access to by virtue of being moved to a separate network. Even if your network configuration is set up as a high-availability environment, or includes intelligent switches that can resolve the virtual machine's needs across different networks, you might experience lag times as the Address Resolution Protocol (ARP) table updates and resumes network traffic for the virtual machines.

Virtual machines reach physical networks through uplink adapters. A vSwitch can transfer data to external networks only when one or more network adapters are attached to it. When two or more adapters are attached to a single vSwitch, they are transparently teamed.

To create or add a virtual network for a virtual machine

1 Log in to the VI Client and select the server from the inventory panel.

The hardware configuration page for this server appears.

2 Click the **Configuration** tab and click **Networking**.

Virtual switches appear in an overview that includes a details layout.

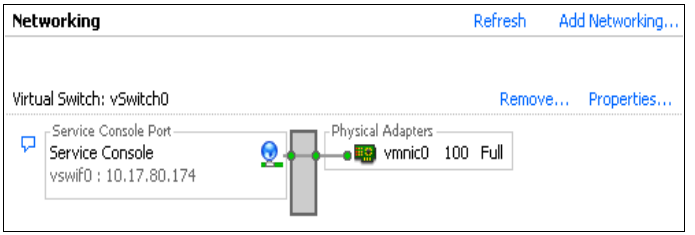

3 On the right side of the page, click **Add Networking**.

NOTE You use the Add Network wizard is to add new ports and port groups.

4 Accept the default connection type, **Virtual Machines**.

Virtual Machines lets you add a labeled network to handle virtual machine network traffic.

5 Click **Next**.

6 Select **Create a virtual switch**.

You can create a new vSwitch with or without Ethernet adapters.

If you create a vSwitch without physical network adapters, all traffic on that vSwitch is confined to that vSwitch. No other hosts on the physical network or virtual machines on other vSwitches can send or receive traffic over this vSwitch. You might create a vSwitch without physical network adapters if you want a group of virtual machines to be able to communicate with each other, but not with other hosts or with virtual machines outside the group.

Changes appear in the **Preview** pane.

7 Click **Next**.

8 In the **Port Group Properties** group, enter a network label that identifies the port group that you are creating.

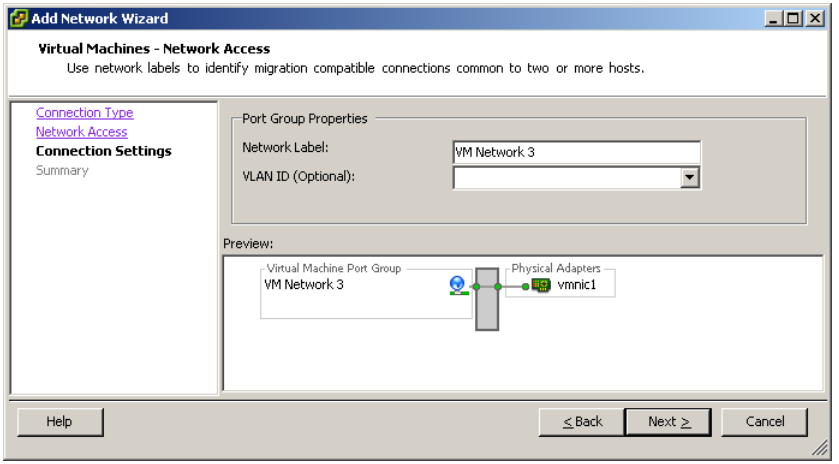

Use network labels to identify migration-compatible connections common to two or more hosts.

9 If you are using a VLAN, in the **VLAN ID** field, enter a number between 1 and 4094.

If you are unsure what to enter, leave this field blank or ask your network administrator.

If you enter 0 or leave the field blank, the port group can see only untagged (non-VLAN) traffic. If you enter 4095, the port group can see traffic on any VLAN while leaving the VLAN tags intact.

10 Click **Next**.

11 After you determine that the vSwitch is configured correctly, click **Finish**.

NOTE To enable failover (NIC teaming), bind two or more adapters to the same switch. If one uplink adapter is not operational, network traffic is routed to another adapter attached to the switch. NIC teaming requires both Ethernet devices to be on the same Ethernet broadcast domain.

VMkernel Networking Configuration

Moving a virtual machine from one host to another is called migration. Migrating a virtual machine that is powered on is called *VMotion*. Migration with VMotion, designed to be used between highly compatible systems, lets you migrate virtual machines with no downtime. Your VMkernel networking stack must be set up properly to accommodate VMotion.

IP Storage refers to any form of storage that uses TCP/IP network communication as its foundation, which includes iSCSI and NFS for ESX Server 3. Because both of these storage types are network based, both types can use the same VMkernel interface and port group.

The network services that the VMkernel (iSCSI, NFS, and VMotion) provides use a TCP/IP stack in the VMkernel. This TCP/IP stack is completely separate from the TCP/IP stack used in the service console. Each of these TCP/IP stacks accesses various networks by attaching to one or more port groups on one or more vSwitches.

TCP/IP Stack at the VMkernel Level

The VMware VMkernel TCP/IP networking stack has been extended to handle iSCSI, NFS, and VMotion in the following ways:

- iSCSI as a virtual machine datastore
- iSCSI for the direct mounting of .ISO files, which are presented as CD-ROMs to virtual machines
- NFS as a virtual machine datastore
- NFS for the direct mounting of `.ISO` files, which are presented as CD-ROMs to virtual machines
- Migration with VMotion

NOTE ESX Server 3 supports only NFS version 3 over TCP/IP.

Implications and Guidelines for Configuration

Refer to the following guidelines when you configure VMkernel networking:

- The IP address that you assign to the service console during installation must be different from the IP address that you assign to the VMkernel's TCP/IP stack from the **Configuration > Networking** tab of the VMware Infrastructure Client.
- Unlike other VMkernel services, iSCSI has a service console component, so networks that are used to reach iSCSI targets must be accessible to both service console and VMkernel TCP/IP stacks.
- Before you configure a software iSCSI for the ESX Server 3 host, open a firewall port by enabling the iSCSI software client service. See "Opening Firewall Ports for Supported Services and Management Agents" on page 188.

To set up the VMkernel

1 Log in to the VMware VI Client and select the server from the inventory panel.

The hardware configuration page for this server appears.

2 Click the **Configuration** tab and click **Networking**.

3 Click **Add Networking**.

4 Select **VMkernel** and click **Next**.

Selecting **VMotion and IP Storage** lets you connect the VMkernel, which runs services for VMotion and IP storage (NFS or iSCSI), to the physical network.

The **Network Access** page appears.

5 Select the vSwitch to use, or select **Create a virtual switch** to create a new vSwitch.

6 Select the check boxes for the network adapters your vSwitch will use.

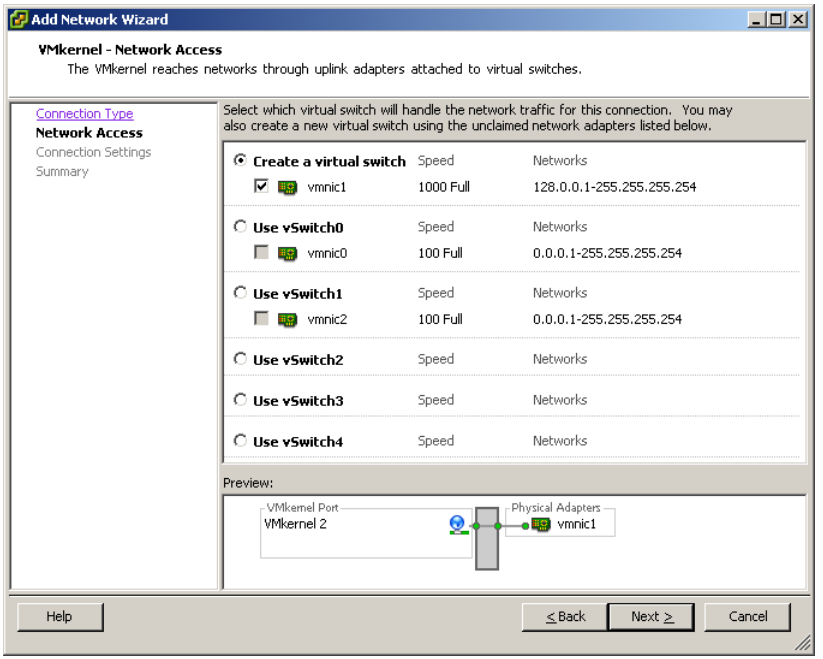

Your choices appear in the **Preview** pane.

Select adapters for each vSwitch so that virtual machines or other services that connect through the adapter can reach the correct Ethernet segment. If no adapters appear under **Create a new virtual switch**, all the network adapters in the system are being used by existing vSwitches. You can either create a new vSwitch without a network adapter, or select a network adapter that an existing vSwitch uses.

For information on moving network adapters between vSwitches, see "To add uplink adapters" on page 45.

7 Click **Next**.

8 In the **Port Group Properties** area, select or enter a network label and a VLAN ID.

- **Network Label** — A name that identifies the port group that you are creating. This is the label that you specify when configuring a virtual adapter to be attached to this port group, when configuring VMkernel services, such as VMotion and IP storage.
- **VLAN ID** — Identifies the VLAN that the port group's network traffic will use.

9 Select **Use this port group for VMotion** to enable this port group to advertise itself to another ESX Server as the network connection where VMotion traffic should be sent.

You can enable this property for only one VMotion and IP storage port group for each ESX Server 3 host. If this property is not enabled for any port group, migration with VMotion to this host is not possible.

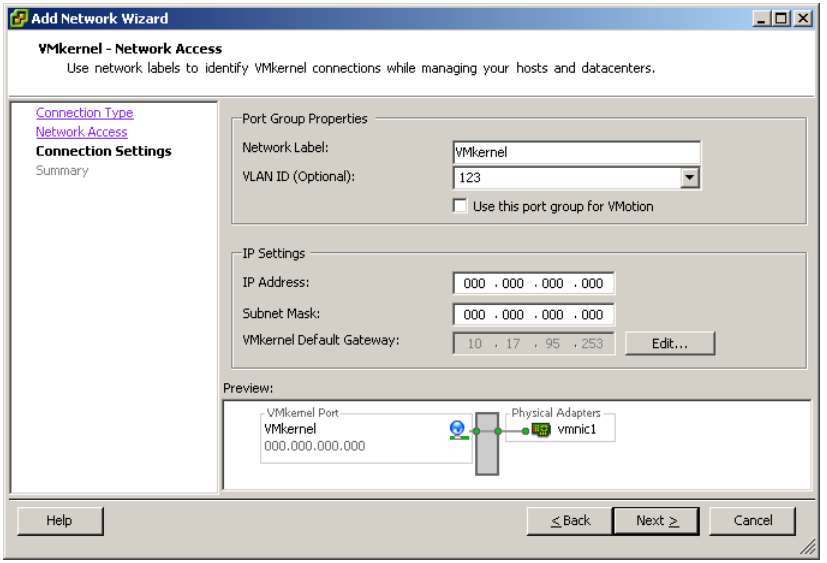

10 In the **IP Settings** group, click **Edit** to set the **VMkernel Default Gateway** for VMkernel services, such as VMotion, NAS, and iSCSI.

NOTE Set a default gateway for the port that you created. VirtualCenter 2 behaves differently from VirtualCenter 1.x. You must use a valid IP address to configure the VMkernel IP stack, not a dummy address.

On the **DNS Configuration** tab, the name of the host is entered into the name field by default. The DNS server addresses that were specified during installation are also preselected, as is the domain.

On the **Routing** tab, the service console and the VMkernel each need their own gateway information. A gateway is needed if connectivity to machines not on the same IP subnet as the service console or VMkernel.

Static IP settings is the default.

11 Click **OK**, then click **Next**.

12 Use the **Back** button to make any changes.

13 Review your changes on the **Ready to Complete** page and click **Finish**.

Service Console Configuration

The service console and the VMkernel use virtual Ethernet adapters to connect to a vSwitch and to reach networks that the vSwitch services.

Basic Service Console Configuration Tasks

Two common service console configuration changes are: changing NICs and changing the settings for an existing NIC that is in use.

When only one service console connection is present, changing the service console configuration is not allowed. For a new connection, change the network settings to use an additional NIC. After you verify that the new connection is functioning properly, remove the old connection. You are switching over to the new NIC.

NOTE You can create a maximum of 16 service console ports in ESX Server 3.

To configure service console networking

1 Log in to the VI Client and select the server from the inventory panel.

The hardware configuration page for this server appears.

2 Click the **Configuration** tab, and click **Networking**.

3 Click **Add Networking**.

4 Select **Service Console** on the **Connection Types** page, and click **Next**.

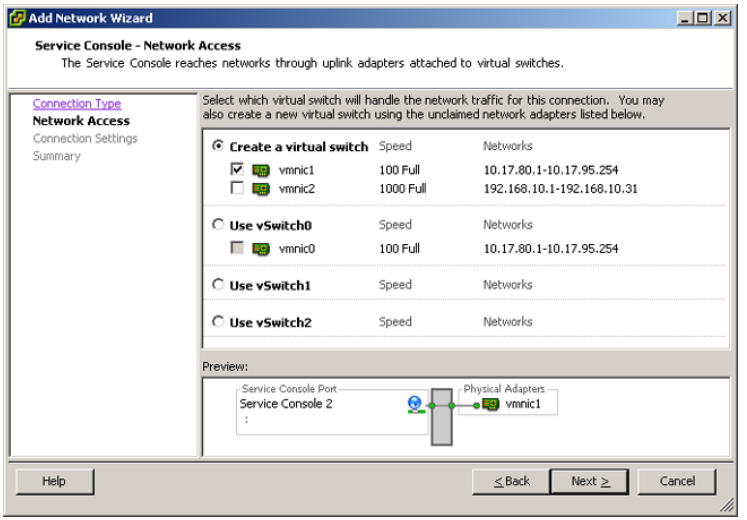

5 Select the vSwitch to use for network access, or select **Create a new vSwitch** and click **Next**.

If no adapters appear in the **Create a new virtual switch** group, all the network adapters in the system are being used by existing vSwitches. For information on moving network adapters between vSwitches, see "To add uplink adapters" on page 45.

6 In the **Port Group Properties** group, select or enter the **Network Label** and **VLAN ID**.

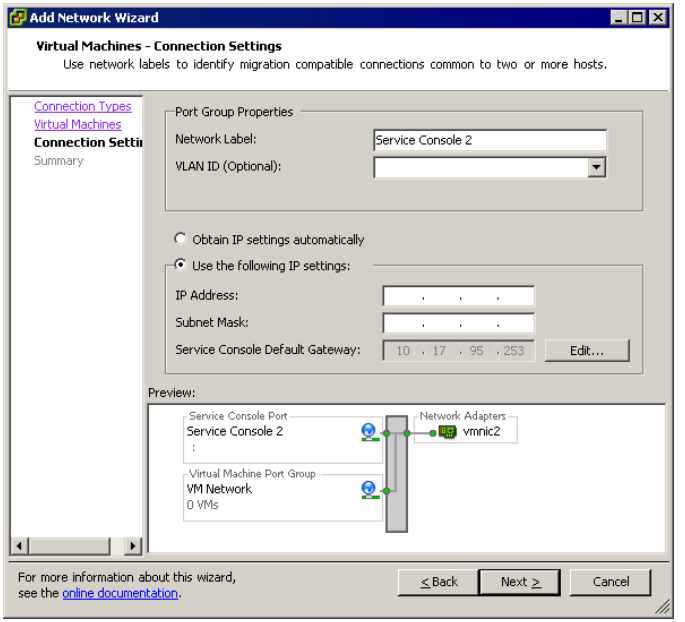

Newer ports and port groups appear at the top of the vSwitch diagram.

7 Enter the **IP Address** and **Subnet Mask,** or select **Obtain IP setting automatically** for the IP address and subnet mask.

8 Click **Edit** to set the **Service Console Default Gateway**.

See "To set the default gateway" on page 37.

9 Click **Next**.

10 Check the information and click **Finish**.

To configure service console ports

1 Log in to the VI Client and select the server from the inventory panel.

The hardware configuration page for this server appears.

2 Click the **Configuration** tab, and click **Networking**.

3 On the right side of the page, find the vSwitch to edit and click **Properties** for that vSwitch.

4 In the **vSwitch Properties** dialog box, click the **Ports** tab.

5 Select S**ervice Console**, and click **Edit**.

A warning dialog box appears to explain that modifying your service console connection might disconnect all management agents.

6 To continue with the service console configuration, click **Continue modifying this connection**.

7 Edit port properties, IP settings, and effective policies as necessary.

8 Click **OK**.

Only one default gateway can be configured per TCP/IP stack.

To set the default gateway

1 Log into the VMware VI Client and select the server from the inventory panel.

The hardware configuration page for this server appears.

2 Click the **Configuration** tab, and click **DNS and Routing**.

The **DNS and Routing** panel appears.

3 Click **Properties**.

On the **DNS Configuration** tab, the name of the host is entered into the name field by default. The DNS server addresses and the domain previously selected during installation are also preselected.

On the **Routing** tab, the service console and the VMkernel are often not connected to the same network, and each needs its own gateway information. A gateway is needed for connectivity to machines not on the same IP subnet as the service console or VMkernel interfaces.

NOTE All NAS and iSCSI servers need to be either reachable by the default gateway, or on the same broadcast domain as the associated vSwitches.

For the service console, the gateway device is needed only when two or more network adapters are using the same subnet. The gateway device determines which network adapter will be used for the default route.

4 Click the **Routing** tab.

5 Set the VMkernel default gateway.

CAUTION There is a risk of misconfiguration, which can cause the UI to lose connectivity to the host, in which case the host will have to be reconfigured from command line at the service console. Be sure that your network settings are correct before saving your changes.

6 Click **OK**.

To display service console information

1 Click the info icon to display service console information.

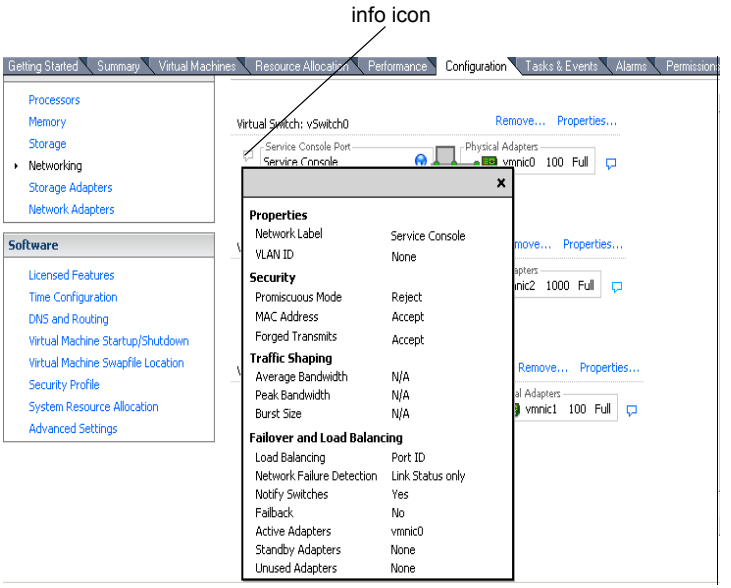

2 Click the **X** to close the information pop-up window.

Using DHCP for the Service Console

In most cases, use static IP addresses for the service console. You can also set up the service console to use dynamic addressing, DHCP, if your DNS server can map the service console's host name to the dynamically generated IP address.

If your DNS server cannot map the host's name to its DHCP-generated IP address, determine the service console's numeric IP address and use that numeric address when accessing the host.

The numeric IP address might change as DHCP leases expire or when the system is rebooted. For this reason, VMware does not recommend using DHCP for the service console unless your DNS server can handle the host name translation.

3 Advanced Networking

This chapter guides you through advanced networking topics in an ESX Server 3 environment, and reviews how to set up and change advanced networking configuration options.

This chapter discusses the following topics:

- "Virtual Switch Properties and Policies" on page 42
- "Port Group Configuration" on page 58
- "DNS and Routing" on page 61
- "TCP Segmentation Offload and Jumbo Frames" on page 62
- "NetQueue and Networking Performance" on page 65
- "Setting Up MAC Addresses" on page 67
- "Networking Best Practices and Tips" on page 69

Virtual Switch Properties and Policies

This section guides you through configuring virtual switch properties and networking policies set at the virtual switch level.

Virtual Switch Properties

Virtual switch settings control vSwitch-wide defaults for ports, which can be overridden by port group settings for each vSwitch.

Editing Virtual Switch Properties

Editing vSwitch properties consists of:

- Configuring ports
- Configuring the uplink network adapters

To edit the number of ports for a vSwitch

1 Log in to the VI Client, and select the server from the inventory panel.

The hardware configuration page for this server appears.

2 Click the **Configuration** tab, and click **Networking**.

3 On the right side of the page, find the vSwitch to edit.

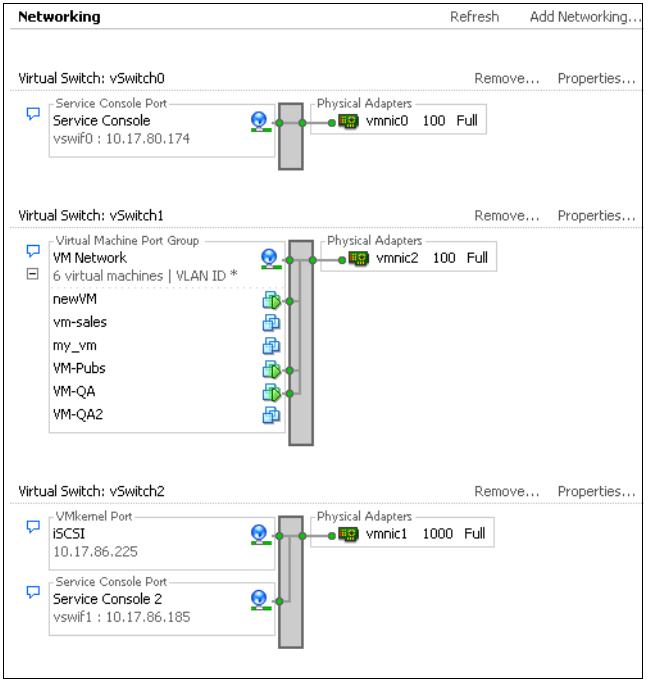

4 Click **Properties** for that vSwitch.

5 Click the **Ports** tab.

6 Select the vSwitch item in the **Configuration** list, and click **Edit**.

7 Click the **General** tab to set the number of ports.

8 Choose the number of ports you want to use from the drop-down menu.

9 Click **OK**.

To configure the uplink network adapter by changing its speed

1 Log in to the VI Client and select the server from the inventory panel.

The hardware configuration page for this server appears.

2 Click the **Configuration** tab, and click **Networking**.

3 Select a vSwitch and click **Properties**.

4 Click the **Network Adapters** tab.

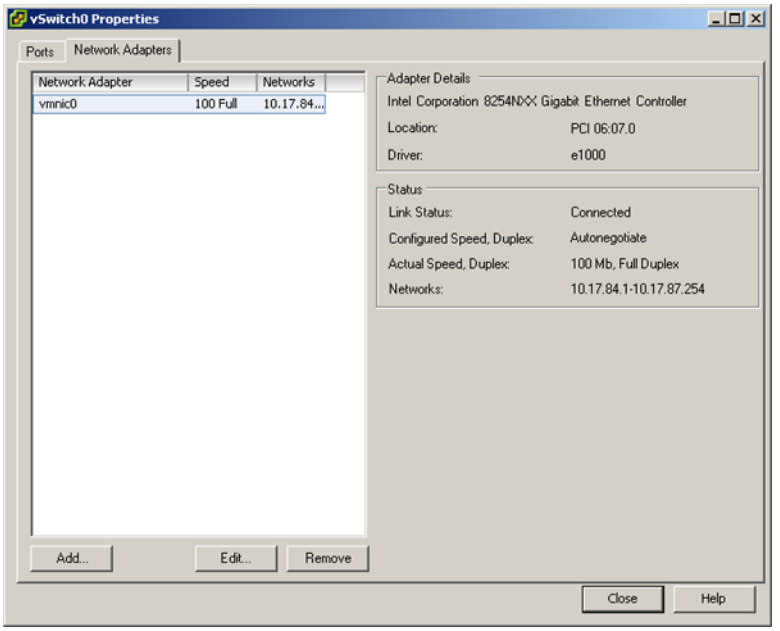

5 To change the configured speed and duplex value of a network adapter, select the network adapter and click **Edit**.

The **Status** dialog box appears. The default is **Autonegotiate**, which is usually the correct choice.

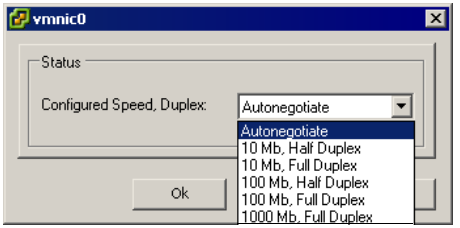

6 To select the connection speed manually, select the speed/duplex from the drop-down menu.

Choose the connection speed manually if the NIC and a physical switch might fail to negotiate the proper connection speed. Symptoms of mismatched speed and duplex include low bandwidth or no link connectivity at all.

The adapter and the physical switch port it is connected to must be set to the same value, that is, auto and auto or ND and ND where ND is some speed and duplex, but not auto and ND.

7 Click **OK**.

To add uplink adapters

1 Log in to the VI Client, and select the server from the inventory panel.

The hardware configuration page for this server appears.

2 Click the **Configuration** tab and click **Networking**.

3 Select a vSwitch and click **Properties**.

4 In the **Properties** dialog box, click the **Network Adapters** tab.

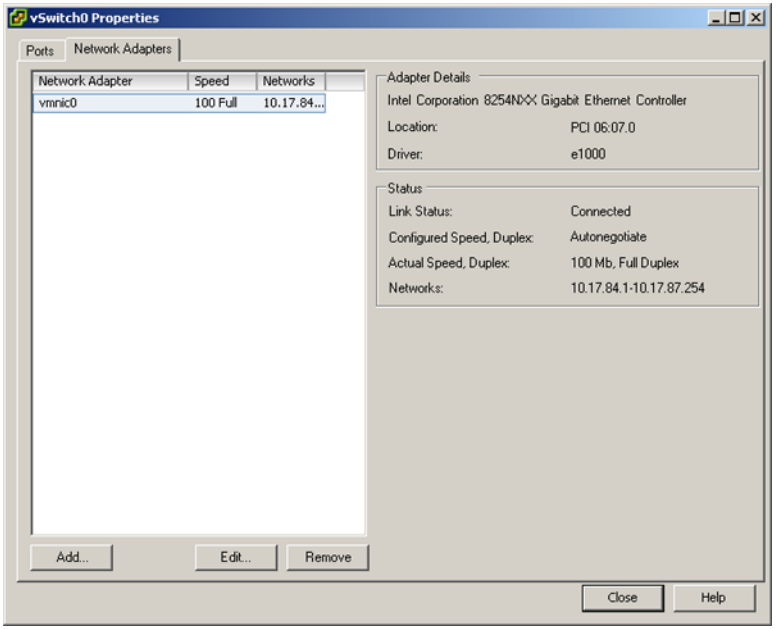

5 Click **Add** to launch the Add Adapter wizard.

You can associate multiple adapters to a single vSwitch to provide NIC teaming. Such a team can share traffic and provide failover.

CAUTION Misconfiguration can result in the loss of the VI Client ability to connect to the host.

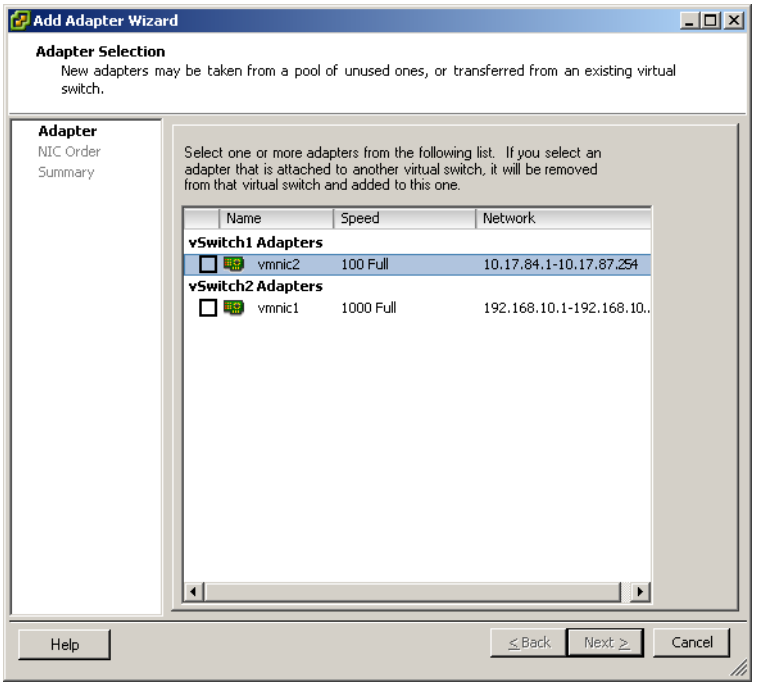

6 Select one or more adapters from the list and click **Next**.

7 To order the NICs, select a NIC and click **Move Up** and **Move Down** to move it up or down into the appropriate category (Active or Standby).

- **Active Adapters** — Adapters that the vSwitch uses.
- **Standby Adapters** — Adapters that become active if one or more of the active adapters fails.

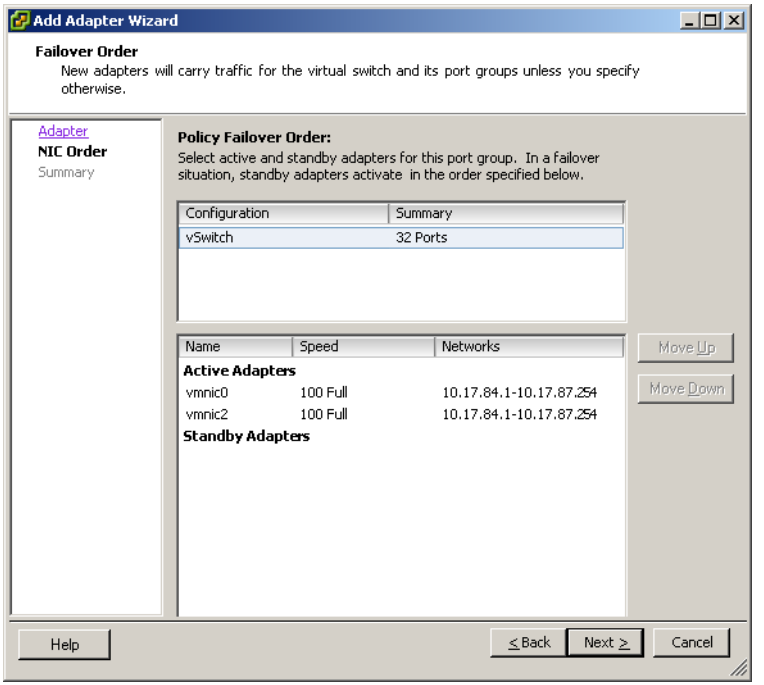

8 Click **Next**.

9 Review the information on the **Adapter Summary** page, click **Back** to change any entries, and click **Finish**.

The list of network adapters reappears, showing the adapters that the vSwitch now claims.

10 Click **Close** to exit the **vSwitch Properties** dialog box.

The **Networking** section in the **Configuration** tab shows the network adapters in their designated order and categories.

Cisco Discovery Protocol

Cisco Discovery Protocol (CDP) allows ESX Server 3 administrators to determine which Cisco switch port is connected to a given vSwitch. When CDP is enabled for a particular vSwitch, you can view properties of the Cisco switch (such as device ID, software version, and timeout) from the VI Client.

You can use the service console command-line interface to enable CDP.

To enable CDP

1 Log in directly to your ESX Server 3 host's console.

2 Use the `esxcfg-vswitch -b <vSwitch>` command to view the current CDP mode for the vSwitch.

If CDP is disabled, the mode will be shown as **down**.

3 Use the `esxcfg-vswitch -B <mode> <vSwitch>` command to change the CDP mode.

The available CDP modes are:

- **down** – CDP is disabled.
- **listen** — ESX Server 3 detects and displays information about the associated Cisco switch port, but information about the vSwitch is not available to the Cisco switch administrator.
- **advertise** – ESX Server 3 makes information about the vSwitch available to the Cisco switch administrator, but does not detect and display information about the Cisco switch
- **both** – ESX Server 3 detects and displays information about the associated Cisco switch and makes information about the vSwitch available to the Cisco switch administrator.

To view Cisco switch information from the VI Client

1 Set the CDP mode for the vSwitch to either **both** or **listen**.

2 Log in to the VI Client, and select the server from the inventory panel.

The hardware configuration page for this server appears.

3 Click the **Configuration** tab, and click **Networking**.

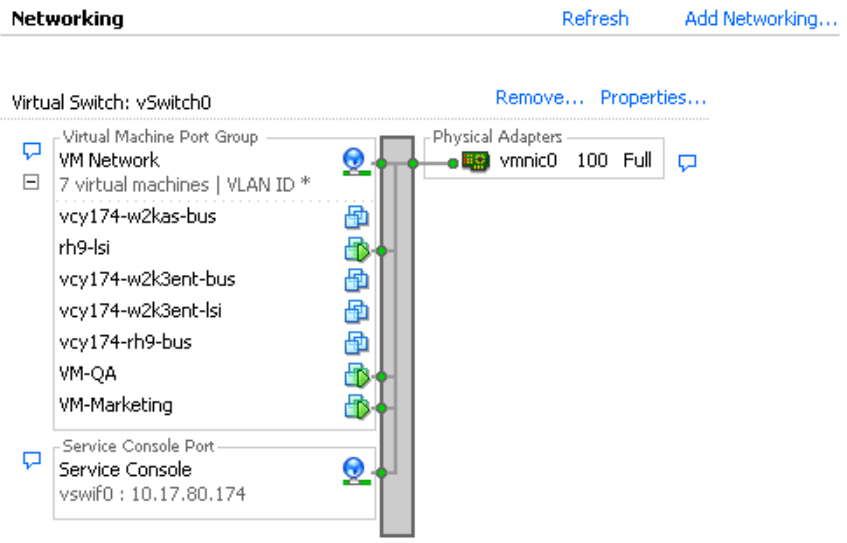

4 Click the info icon to the right of the vSwitch.

Cisco Discovery Protocol

Properties	
Version	0
Timeout	0
Time to live	142
Samples	22
Device Id	blade-vlan-switch
Address	28.20.17.10
Port Id	GigabitEthernet0/8
Software Version	Cisco Internetwork Operati
Hardware Platform	cisco WS-C2970G-24T-E
IP Prefix	0.0.0.0
IP Prefix Length	0
VLAN	1
Full Duplex	true
MTU	0
System Name	
System OId	
Management Address	28.20.17.10
Location	
CDP Device Capability	
Router	false
Transparent Bridge	false
Source Route Bridge	false
Network Switch	true
Host	false
IGMP Enabled	true
Repeater	false

NOTE Because the CDP advertisements of Cisco equipment typically occur once a minute, a noticeable delay might occur between enabling CDP on ESX Server 3 and the availability of CDP data from the VI client.

Virtual Switch Policies

You can apply a set of vSwitch-wide policies by selecting the vSwitch at the top of the **Ports** tab and clicking **Edit**.

To override any of these settings for a port group, select that port group and click **Edit**. Any changes to the vSwitch-wide configuration are applied to any of the port groups on that vSwitch, except for the configuration options that are overridden by the port group.

The vSwitch policies consist of:

- Layer 2 Security policy
- Traffic Shaping policy
- Load Balancing and Failover policy

Layer 2 Security Policy

Layer 2 is the data link layer. The three elements of the Layer 2 Security policy are promiscuous mode, MAC address changes, and forged transmits.

In nonpromiscuous mode, a guest adapter listens to traffic only on its own MAC address. In promiscuous mode, it can listen to all the packets. By default, guest adapters are set to non-promiscuous mode.

For further information on security, see "Securing Virtual Switch Ports" on page 202.

To edit the Layer 2 Security policy

1 Log into the VMware VI Client and select the server from the inventory panel.

The hardware configuration page for this server appears.

2 Click the **Configuration** tab, and click **Networking**.

3 Click **Properties** for the vSwitch whose Layer 2 Security policy you want to edit.

4 In the **Properties** dialog box for the vSwitch, click the **Ports** tab.

5 Select the vSwitch item and click **Edit**.

6 In the Properties dialog box for the vSwitch, click the **Security** tab.

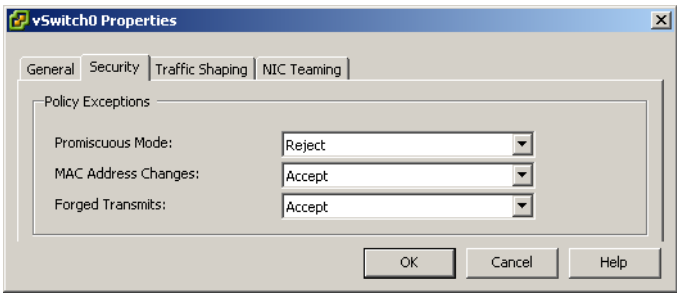

By default, **Promiscuous Mode** is set to **Reject**, **MAC Address Changes**, and **Forced Transmits** are set to **Accept**.

The policy here applies to all virtual adapters on the vSwitch except where the port group for the virtual adapter specifies a policy exception.

7 In the **Policy Exceptions** pane, select whether to reject or accept the Layer2 Security policy exceptions:

- **Promiscuous Mode**
  - **Reject** — Placing a guest adapter in promiscuous mode has no effect on which frames are received by the adapter.
  - **Accept** — Placing a guest adapter in promiscuous mode causes it to detect all frames passed on the vSwitch that are allowed under the VLAN policy for the port group that the adapter is connected to.
- **MAC Address Changes**
  - **Reject** — If you set the **MAC Address Changes** to **Reject** and the guest operating system changes the MAC address of the adapter to anything other than what is in the `.vmx` configuration file, all inbound frames are dropped.

    If the Guest OS changes the MAC address back to match the MAC address in the `.vmx` configuration file, inbound frames are passed again.
  - **Accept** — Changing the MAC address from the Guest OS has the intended effect: frames to the new MAC address are received.
- **Forged Transmits**
  - **Reject** — Any outbound frame with a source MAC address that is different from the one currently set on the adapter are dropped.
  - **Accept** — No filtering is performed and all outbound frames are passed.

8 Click **OK**.

Traffic Shaping Policy

ESX Server 3 shapes traffic by establishing parameters for three outbound traffic characteristics: average bandwidth, burst size, and peak bandwidth. You can set values for these characteristics through the VI Client, establishing a traffic shaping policy for each port group.

- **Average Bandwidth** establishes the number of bits per second to allow across the vSwitch averaged over time—the allowed average load.
- **Burst Size** establishes the maximum number of bytes to allow in a burst. If a burst exceeds the burst size parameter, excess packets are queued for later transmission. If the queue is full, the packets are dropped. When you specify values for these two characteristics, you indicate what you expect the vSwitch to handle during normal operation.
- **Peak Bandwidth** is the maximum bandwidth the vSwitch can absorb without dropping packets. If traffic exceeds the peak bandwidth that you establish, excess packets are queued for later transmission after traffic on the connection returns to the average and enough spare cycles are available to handle the queued packets. If the queue is full, the packets are dropped. Even if you have spare bandwidth because the connection has been idle, the peak bandwidth parameter limits transmission to no more than peak until traffic returns to the allowed average load.

To edit the Traffic Shaping policy

1 Log in to the VI Client and select the server from the inventory panel.

  The hardware configuration page for this server appears.

2 Click the **Configuration** tab, and click **Networking**.

3 Select a vSwitch and click **Properties**.

4 In the **vSwitch Properties** dialog box, click the **Ports** tab.

5 Select the vSwitch and click **Edit**.

6 Click the **Traffic Shaping** tab.

When traffic shaping is disabled, the tunable features are dimmed. You can selectively override all traffic-shaping features at the port group level if traffic shaping is enabled.

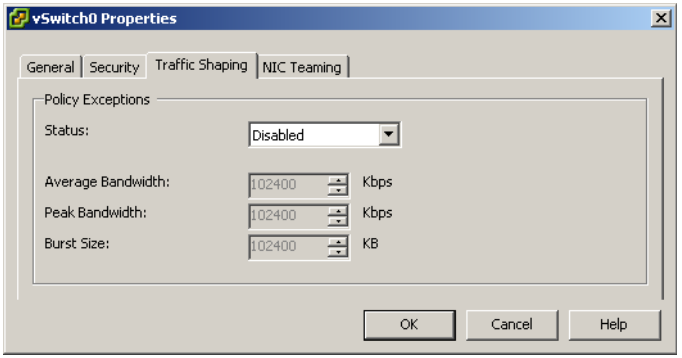

This policy is applied to each individual virtual adapter attached to the port group, not to the vSwitch as a whole.

Status — If you enable the policy exception in the **Status** field, you are setting limits on the amount of networking bandwidth allocation for each virtual adapter associated with this particular port group. If you disable the policy, services have a free, clear connection to the physical network by default.

The remaining fields define network traffic parameters:

- **Average Bandwidth** is a value measured over a particular period of time.
- **Peak Bandwidth** is a value that is the maximum bandwidth allowed and that can never be smaller than average bandwidth. This parameter limits the maximum bandwidth during a burst.
- **Burst Size** is a value that specifies how large a burst can be in kilobytes (KB). This parameter controls the amount of data that can be sent in one burst.

Load Balancing and Failover Policy

Load Balancing and Failover policies allow you to determine how network traffic is distributed between adapters and how to re-route traffic in the event of an adapter failure by configuring the following parameters:

- **Load Balancing policy** determines how outgoing traffic is distributed among the network adapters assigned to a vSwitch.

  NOTE Incoming traffic is controlled by the Load Balancing policy on the physical switch.

- **Failover Detection**: Link Status and Beacon Probing
- **Network Adapter Order** (Active or Standby)

To edit the failover and load balancing policy

1. Log in to the VI Client and select the server from the inventory panel.

   The hardware configuration page for this server appears.

2. Click the **Configuration** tab, and click **Networking**.
3. Select a vSwitch and click **Edit**.
4. In the **vSwitch Properties** dialog box, click the **Ports** tab.
5. To edit the **Failover and Load Balancing** values for the vSwitch, select the vSwitch item and click **Properties**.

6 Click the **NIC Teaming** tab.

You can override the failover order at the port group level. By default, new adapters are active for all policies. New adapters carry traffic for the vSwitch and its port group unless you specify otherwise.

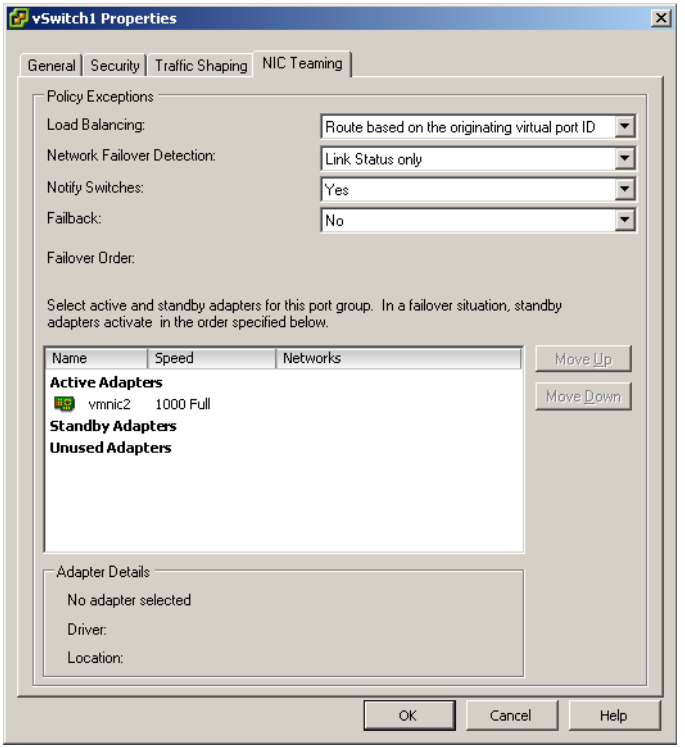

7 In the **Policy Exceptions** group:

- **Load Balancing** — Specify how to choose an uplink.
  - **Route based on the originating port ID** — Choose an uplink based on the virtual port where the traffic entered the virtual switch.
  - **Route based on ip hash** — Choose an uplink based on a hash of the source and destination IP addresses of each packet. For non-IP packets, whatever is at those offsets is used to compute the hash.

- **Route based on source MAC hash** — Choose an uplink based on a hash of the source Ethernet.
- **Use explicit failover order** — Always use the highest order uplink from the list of Active adapters which passes failover detection criteria.

NOTE IP-based teaming requires that the physical switch be configured with etherchannel. For all other options, etherchannel should be disabled.

- **Network Failover Detection** — Specify the method to use for failover detection.
  - **Link Status only** – Relies solely on the link status that the network adapter provides. This option detects failures, such as cable pulls and physical switch power failures, but not configuration errors, such as a physical switch port being blocked by spanning tree or that is misconfigured to the wrong VLAN or cable pulls on the other side of a physical switch.
  - **Beacon Probing** – Sends out and listens for beacon probes on all NICs in the team and uses this information, in addition to link status, to determine link failure. This detects many of the failures previously mentioned that are not detected by link status alone.
- **Notify Switches** — Select **Yes** or **No** to notify switches in the case of failover.

  If you select **Yes**, whenever a virtual NIC is connected to the vSwitch or whenever that virtual NIC's traffic would be routed over a different physical NIC in the team because of a failover event, a notification is sent out over the network to update the lookup tables on physical switches. In almost all cases, this process is desirable for the lowest latency of failover occurrences and migrations with VMotion.

NOTE Do not use this option when the virtual machines using the port group are using Microsoft Network Load Balancing in unicast mode. No such issue exists with NLB running in multicast mode.

- **Failback** — Select **Yes** or **No** to disable or enable failback.

  This option determines how a physical adapter is returned to active duty after recovering from a failure. If failback is set to **Yes** (default), the adapter is returned to active duty immediately upon recovery, displacing the standby adapter that took over its slot, if any. If failback is set to **No**, a failed adapter is left inactive even after recovery until another currently active adapter fails, requiring its replacement.

- **Failover Order** — Specify how to distribute the work load for adapters. If you want to use some adapters but reserve others for emergencies in case the adapters in use fail, set this condition by using the drop-down menu to place them into the two groups:
  - **Active Adapters** — Continue to use the adapter when the network adapter connectivity is up and active.
  - **Standby Adapters** — Use this adapter if one of the active adapter's connectivity is down.
  - **Unused Adapters** — Not to be used.

Port Group Configuration

You can change the following port group configurations:

- Port group properties
- Labelled network policies

To edit port group properties

1 Log into the VMware VI Client, and select the server from the inventory panel.

  The hardware configuration page for this server appears.

2 Click the **Configuration** tab, and click **Networking**.

3 On the right side of the window, click **Properties** for a network.

4 Click the **Ports** tab.

5 Select the port group and click **Edit**.

6 In the **Properties** dialog box for the port group, click the **General** tab to change:

  - **Network Label** — Identifies the port group that you are creating. Specify this label when you configure a virtual adapter to be attached to this port group, either when you configure virtual machines or when you configure VMkernel services, such as VMotion and IP storage.
  - **VLAN ID** — Identifies the VLAN that the port group's network traffic will use.

7 Click **OK**.

To override labeled network policies

1 To override any network policies for a particular labeled network, select the network, click **Edit**, and click the **Security** tab.

2 Select the check box for the labeled network policy to override.

For information on these settings, see "Layer 2 Security Policy" on page 51.

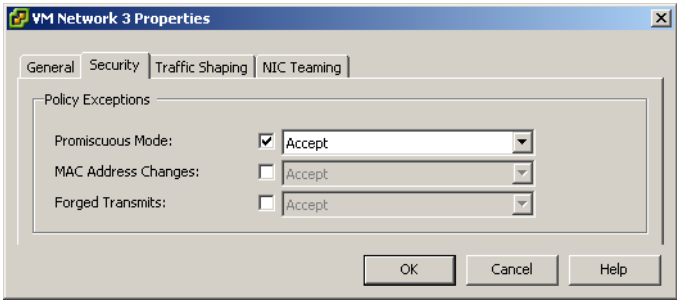

3 Click **the Traffic Shaping tab**.

4 Select the check box next to **Status** and select **Enabled** or **Disabled**.

For information on the **Status** settings, see "Traffic Shaping Policy" on page 53.

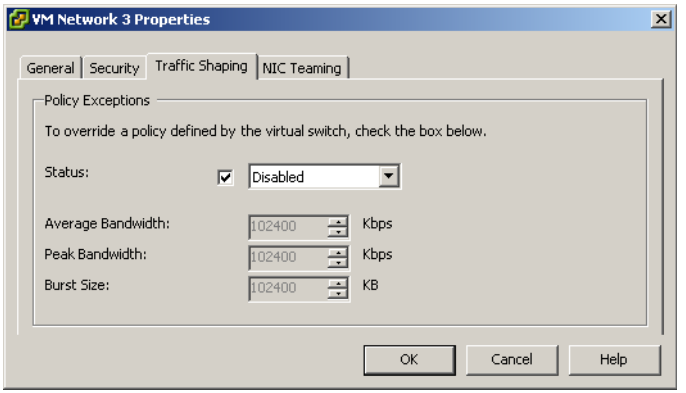

5 Click the **NIC Teaming** tab.

6 Select the associated check box to override the load balancing or failover order policies.

For information on these settings, see "Load Balancing and Failover Policy" on page 55.

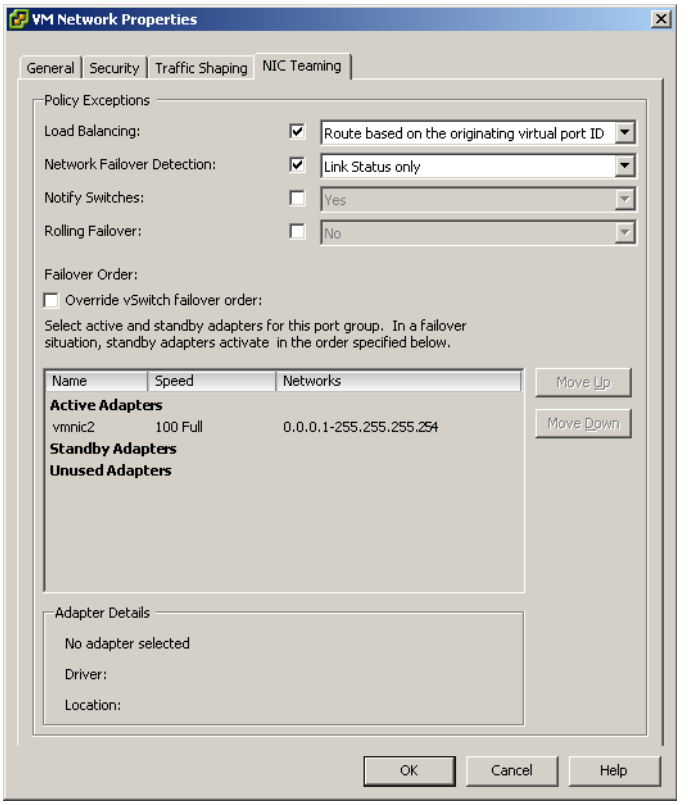

7 Click **OK**.

DNS and Routing

Configure DNS and routing through the VI Client.

To change the DNS and routing configuration

1 Log in to the VI Client and select the server from the inventory panel.

The hardware configuration page for this server appears.

2 Click the **Configuration** tab, and click **DNS and Routing**.

3 On the right of the window, click **Properties**.

4 In the **DNS Configuration** tab, enter values for the **Name** and **Domain** fields.

5 Choose to obtain the DNS server address automatically or use a DNS server address.

NOTE DHCP is supported only if the DHCP server is accessible to the service console. In other words, the service console must have a virtual interface (vswif) configured and attached to the network where the DHCP server resides.

---

6 Specify the domains in which to look for hosts.

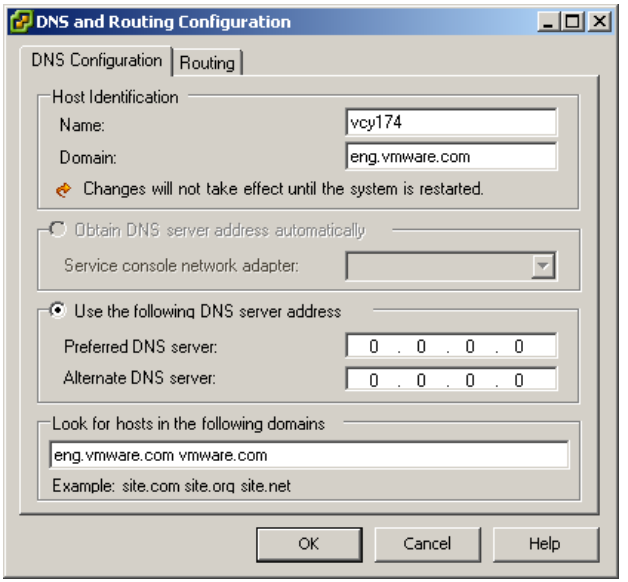

7 On the **Routing** tab, change default gateway information as needed.

Select a gateway device only if you have configured the service console to connect to more than one subnet.

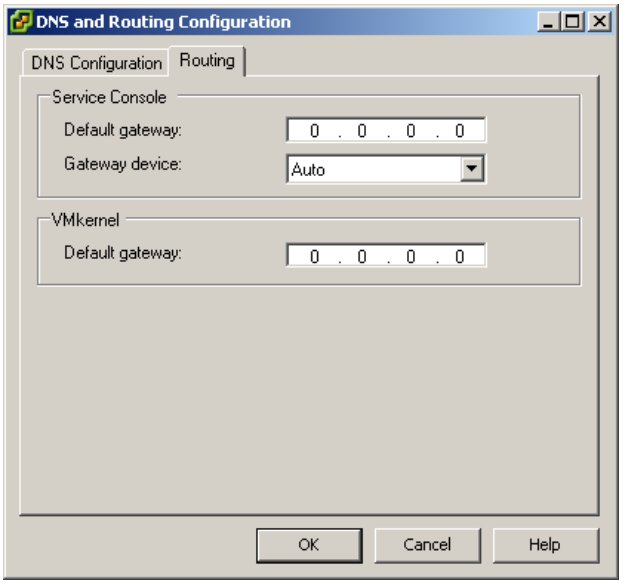

8 Click **OK**.

TCP Segmentation Offload and Jumbo Frames

TCP Segmentation Offload (TSO) and Jumbo Frame support are added to the TCP/IP stack in ESX Server 3 version 3.5. Jumbo Frames must be enabled at the server level using the command-line interface to configure the MTU size for each vSwitch. TSO is enabled on the VMkernel interface by default, but must be enabled at the virtual machine level.

Enabling TSO

TSO support through the Enhanced vmxnet network adapter is available for virtual machines running the following guest operating systems:

- Microsoft Windows 2003 Enterprise Edition with Service Pack 2 (32-bit and 64-bit)
- Red Hat Enterprise Linux 4 (64-bit)

- Red Hat Enterprise Linux 5 (32-bit and 64-bit)
- SuSE Linux Enterprise Server 10 (32-bit and 64-bit)

To enable TSO at the virtual machine level, you must replace the existing vmxnet or Flexible virtual network adapters with Enhanced vmxnet virtual network adapters. This may result in a change in the MAC address of the virtual network adapter.

To enable TSO support for a virtual machine

1 Log in to the VI Client and select the virtual machine from the inventory panel.

   The hardware configuration page for this server appears.

2 Click the **Summary** tab, and click **Edit Settings**.

3 Select the network adapter from the **Hardware** list.

4 Record the network settings and MAC address that the network adapter is using.

5 Click **Remove** to remove the network adapter from the virtual machine.

6 Click **Add**.

7 Select **Ethernet Adapter** and click **Next**.

8 In the **Adapter Type** group, select **Enhanced vmxnet**.

9 Select the network setting and MAC address that the old network adapter was using and click **Next**.

10 Click **Finish**.

11 Click **OK**.

12 If the virtual machine is not set to upgrade VMware Tools at each power-on, you must upgrade VMware Tools manually. See the *Basic System Administration Guide*.

TSO is enabled by default on a VMkernel interface. If TSO gets disabled for a particular VMkernel interface, the only way to enable TSO is to delete that VMkernel interface and re-create it with TSO enabled.

To check that TSO is enabled on a VMkernel interface

1 Log in directly to your ESX Server 3 host's console.

2 Use the `esxcfg-vmknic -l` command to display a list of VMkernel interfaces.

   Each TSO-enabled VMkernel interface should appear on the list with **TSO MSS** set to 40960.

If TSO is not enabled for a particular VMkernel interface, the only way to enable TSO is to delete that VMkernel interface and re-create that VMkernel interface. See "VMkernel Networking Configuration" on page 30.

Enabling Jumbo Frames

Jumbo Frames allow ESX Server 3 to send larger frames out onto the physical network. The network must support Jumbo Frames end-to-end for Jumbo Frames to be effective. Jumbo Frames up to 9kB (9000 Bytes) are supported. iSCSI with Jumbo Frames is not supported.

Jumbo Frames must be enabled for each vSwitch or VMkernel interface through the command-line interface on your ESX Server 3 host. Before enabling Jumbo Frames, check with your hardware vendor to ensure your physical network adapter supports Jumbo Frames.

To create a Jumbo Frames-enabled vSwitch

1 Log in directly to your ESX Server 3 host's console.

2 Use the `esxcfg-vswitch -m <MTU> <vSwitch>` command to set the MTU size for the vSwitch.

This command sets the MTU for all uplinks on that vSwitch. The MTU size should be set to the largest MTU size among all the virtual network adapters connected to the vSwitch.

3 Use the `esxcfg-vswitch -l` command to display a list of vSwitches on the host, and check that the configuration of the vSwitch is correct.

To enable Jumbo Frame support on a virtual machine

1 Log in to the VI Client and select the virtual machine from the inventory panel.

The hardware configuration page for this server appears.

2 Click the **Summary** tab, and click **Edit Settings**.

3 Select the network adapter from the **Hardware** list.

4 Record the network settings and MAC address that the network adapter is using.

5 Click **Remove** to remove the network adapter from the virtual machine.

6 Click **Add**.

7 Select **Ethernet Adapter** and click **Next**.

8 In the **Adapter Type** group, select **Enhanced vmxnet**.

9 Select the network setting and MAC address that the old network adapter was using and click **Next**.

10 Click **Finish**.

11 Select the new network adapter from the **Hardware** list.

12 Under **MAC Address**, select **Manual**, and enter the MAC address that the old network adapter was using.

13 Click **OK**.

14 Check that the Enhanced vmxnet adapter is connected to a vSwitch with Jumbo Frames enabled.

15 Inside the guest operating system, configure the network adapter to allow Jumbo Frames. See your guest operating system's documentation for details.

16 Configure all physical switches and any physical or virtual machines to which this virtual machine connects to support Jumbo Frames.

Additionally, the procedure for enabling a VMkernel interface is incomplete. Please use the following procedure.

To create a Jumbo Frames-enabled VMkernel interface

1 Log in directly to your ESX Server 3 host's console

2 Use the `esxcfg-vmknic -a -i <ip address> -n <netmask> -m <MTU> <port group name>` command to create a VMkernel connection with Jumbo Frame support.

3 Use the `esxcfg-vmknic -l` command to display a list of VMkernel interfaces, and check that the configuration of the Jumbo Frame-enabled interface is correct.

4 Check that the VMkernel interface is connected to a vSwitch with Jumbo Frames enabled.

5 Configure all physical switches and any physical or virtual machines to which this VMkernel interface connects to support Jumbo Frames.

NetQueue and Networking Performance

NetQueue in ESX Server 3 takes advantage of the capability of some network adapters to deliver network traffic to the system in multiple receive queues that can be processed separately. This allows processing to be scaled to multiple CPUs, improving receive-size networking performance.

To enable NetQueue on an ESX Server 3 host

1 Log in to the VI Client and select the server from the inventory panel.

The hardware configuration page for this server appears.

2 Click the **Configuration** tab, and click **Advanced Settings**.

3 Select **VMkernel**.

4 Select **VMkernel.Boot.netNetQueueEnable** and click **OK**.

5 Log in directly to your ESX Server 3 host's console to configure your NIC driver to use NetQueue.

- If you are using the s2io NIC driver, use the `esxcfg-module -s "intr_type=2 rx_ring_num=8" s2io` command to set the appropriate parameters on the s2io module.
- If you are using the ixgbe NIC driver, use the `esxcfg-module -s "InterruptType=2 MQ=1 VMDQ=16" ixgbe` command to set the appropriate parameters on the ixgbe module.
- For third party drivers, contact the third party vendor for appropriate configurations.

6 Reboot the ESX Server 3 host.

To disable NetQueue options on an ESX Server 3 host

1 Log in to the VI Client and select the server from the inventory panel.

The hardware configuration page for this server appears.

2 Click the **Configuration** tab, and click **Advanced Settings**.

3 Deselect **VMkernel.Boot.netNetQueueEnable** and click **OK**.

4 Log in directly to your ESX Server 3 host's console.

5 Use the `esxcfg-module -s "" [module name]` command.

For example, if you are using the s2io NIC driver, use `esxcfg-module -s "" s2io`

6 Reboot the ESX Server 3 host.

Setting Up MAC Addresses

MAC addresses are generated for virtual network adapters that the service console, the VMkernel and virtual machines use. In most cases, these MAC addresses are appropriate. However, you might need to set a MAC address for a virtual network adapter, as in the following cases:

- Virtual network adapters on different physical servers share the same subnet and are assigned the same MAC address, causing a conflict.
- To ensure that a virtual network adapter always has the same MAC address.

The following sections describe how MAC addresses are generated and how you can set the MAC address for a virtual network adapter.

MAC Addresses Generation

Each virtual network adapter in a virtual machine is assigned its own unique MAC address. A MAC address is a six-byte number. Each network adapter manufacturer is assigned a unique three-byte prefix called an OUI (Organizationally Unique Identifier) that it can use to generate unique MAC addresses.

VMware has the following OUIs:

- One for generated MAC addresses.
- One for manually set MAC addresses.
- One that was previously used for legacy virtual machines, but is no longer used with ESX Server 3.

The first three bytes of the MAC address that is generated for each virtual network adapter are comprised of the OUI. This MAC address-generation algorithm produces the other three bytes. The algorithm guarantees unique MAC addresses within a machine and attempts to provide unique MAC addresses across machines.

The network adapters for each virtual machine on the same subnet should have unique MAC addresses. Otherwise, they can behave unpredictably. The algorithm puts a limit on the number of running and suspended virtual machines at any one time on any given server. It also does not handle all cases when virtual machines on distinct physical machines share a subnet.

The VMware Universally Unique Identifier (UUID) generates MAC addresses that are checked for any conflicts. The generated MAC addresses are created by using three parts: the VMware OUI, the SMBIOS UUID for the physical ESX Server 3 machine, and a hash based on the name of the entity that the MAC address is being generated for.

After the MAC address is generated, it does not change unless the virtual machine is moved to a different location, for example, to a different path on the same server. The MAC address in the configuration file of the virtual machine is saved. All MAC addresses that have been assigned to network adapters of running and suspended virtual machines on a given physical powered-off virtual machine is not checked against those of running or suspended virtual machines. It is possible but unlikely that when a virtual machine is powered on again, it can acquire a different MAC address. This acquisition is due to a conflict with a virtual machine that was powered on when this virtual machine was powered off.

Setting MAC Addresses

To circumvent the limit of 256 virtual network adapters per physical machine and possible MAC address conflicts between virtual machines, system administrators can manually assign MAC addresses. VMware uses this OUI for manually generated addresses: 00:50:56.

The MAC address range is

```
00:50:56:00:00:00–00:50:56:3F:FF:FF
```

You can set the addresses by adding the following line to a virtual machine's configuration file:

```
ethernet <number>.address = 00:50:56:XX:YY:ZZ
```

where <number> refers to the number of the Ethernet adapter, XX is a valid hexadecimal number between 00 and 3F, and YY and ZZ are valid hexadecimal numbers between 00 and FF. The value for XX must not be greater than 3F to avoid conflict with MAC addresses that are generated by the VMware Workstation and VMware GSX Server products. The maximum value for a manually generated MAC address is

```
ethernet<number>.address = 00:50:56:3F:FF:FF
```

You must also set the option in a virtual machine's configuration file:

```
ethernet<number>.addressType="static"
```

Because VMware ESX Server 3 virtual machines do not support arbitrary MAC addresses, the above format must be used. As long as you choose a unique value for XX:YY:ZZ among your hard-coded addresses, conflicts between the automatically assigned MAC addresses and the manually assigned ones should never occur.

Using MAC Addresses

You can change a powered-down virtual machine's virtual NICs to use statically assigned MAC addresses using the VI Client.

To set up a MAC address

1 Log in to the VI Client and select the virtual machine from the inventory panel.

2 Click the **Summary** tab, and click **Edit Settings**.

3 Select the network adapter from the **Hardware** list.

4 In the **MAC Address** group, select **Manual**.

5 Enter the desired static MAC address, and click **OK**.

Networking Best Practices and Tips

This section provides information about:

- Networking best practices
- Networking tips

Networking Best Practices

Consider these best practices for configuring your network:

- Separate network services from one another to achieve greater security or better performance.

  To have a particular set of virtual machines function at the highest performance levels, put them on a separate physical NIC. This separation allows for a portion of the total networking workload to be more evenly shared across multiple CPUs. The isolated virtual machines can then better serve traffic from a Web client, for instance.

- You can satisfy the following recommendations either by using VLANs to segment a single physical network, or by using separate physical networks (the latter is preferable).
  - Keeping the service console on its own network is an important part of securing the ESX Server 3 system. Consider the service console network connectivity in the same light as any remote access device in a server, because compromise of the service console gives an attacker full control of all virtual machines running on the system.
  - Keeping the VMotion connection on a separate network devoted to VMotion is important because when migration with VMotion occurs, the contents of the guest operating system's memory are transmitted over the network.

Mounting NFS Volumes

In ESX Server 3, the model of how ESX Server 3 accesses NFS storage of ISO images that are used as virtual CD-ROMs for virtual machines is different from the model used in ESX Server 2.x.

ESX Server 3 has support for VMkernel-based NFS mounts. The new model is to mount your NFS volume with the ISO images through the VMkernel NFS functionality. All NFS volumes mounted in this way appear as datastores in the VI Client. The virtual machine configuration editor allows you to browse the service console file system for ISO images to be used as virtual CD-ROM devices.

Networking Tips

Consider the following networking tips:

- To physically separate network services and to dedicate a particular set of NICs to a specific network service, create a vSwitch for each service. If this is not possible, separate them from each other on a single vSwitch by attaching them to port groups with different VLAN IDs. In either case, confirm with your network administrator that the networks or VLANs you choose are isolated in the rest of your environment, that is, no routers connect them.
- You can add and remove NICs from the vSwitch without affecting the virtual machines or the network service that is running behind that vSwitch. If you remove all of the running hardware, the virtual machines can still communicate amongst themselves. Moreover, if you leave one NIC intact, all of the virtual machines can still connect with the physical network.

- To separate virtual machines into groups, use port groups with different sets of active adapters in their teaming policy. These can use separate adapters as long as all adapters are up but still fall back to sharing in the event of a network or hardware failure.
- To protect your most sensitive virtual machines, deploy firewalls in virtual machines that route between virtual networks with uplinks to physical networks and pure virtual networks with no uplinks.

Networking Scenarios and Troubleshooting 4

This chapter describes common networking configuration and troubleshooting scenarios.

This chapter discusses the following topics:

- "Networking Configuration for Software iSCSI Storage" on page 73
- "Configuring Networking on Blade Servers" on page 78
- "Troubleshooting" on page 81

Networking Configuration for Software iSCSI Storage

The storage you configure for an ESX Server 3 host might include one or more storage area networks (SANs) that use iSCSI storage, which is a means of accessing SCSI devices and exchanging data records by using TCP/IP protocol over a network port rather than through a direct connection to a SCSI device. In iSCSI transactions, blocks of raw SCSI data are encapsulated in iSCSI records and transmitted to the requesting device or user.

NOTE Software-initiated iSCSI is not available over 10GigE network adapters in ESX Server 3.5.

Before you can configure iSCSI storage, you must create a VMkernel port to handle iSCSI networking and a service console connection to the iSCSI network.

To create a VMkernel port for software iSCSI

1 Log in to the VI Client and select the server from the inventory panel.

The hardware configuration page for this server appears.

2 Click the **Configuration** tab, and click **Networking**.

3 Click **Add Networking**.

4 Select **VMkernel** and click **Next**.

The **Network Access** page appears. On this page you connect the physical network to the VMkernel, which runs services for iSCSI storage.

5 Select the vSwitch to use or click **Create a virtual switch**.

6 Select the check boxes for the network adapters that your vSwitch will use.

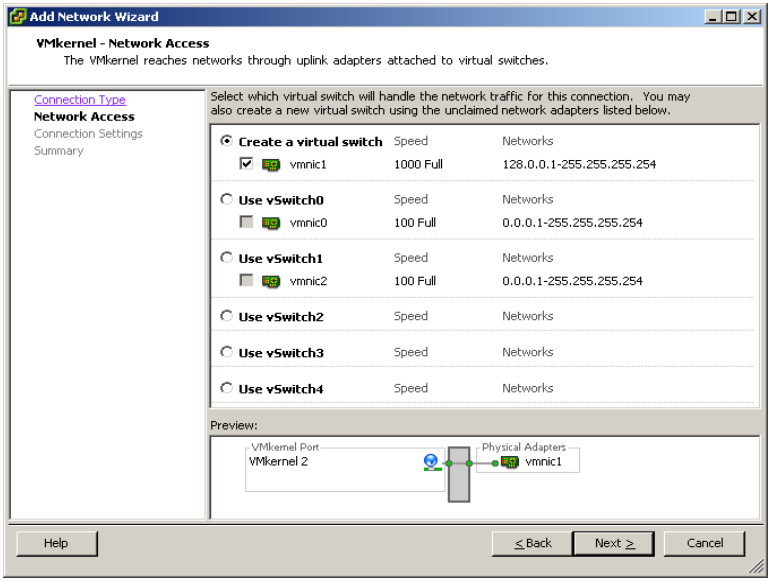

Your choices appear in the **Preview** pane.

Select adapters for each vSwitch so that virtual machines or other services that connect through the adapter can reach the correct Ethernet segment. If no adapters appear in the **Create a virtual switch** group, all of the network adapters in the system are being used by existing vSwitches.

For information on moving network adapters between vSwitches, see "To add uplink adapters" on page 45.

NOTE Do not use iSCSI on 100MB network adapters.

7 Click **Next**.

8 In the **Port Group Properties** group, select or enter a network label and, optionally, a VLAN ID.

Network Label. A name that identifies the port group that you are creating. When you configure iSCSI storage, specify this label when you configure a virtual adapter to be attached to this port group.

VLAN ID. Identifies the VLAN that the port group's network traffic will use. VLAN IDs are not required. If you are not sure whether or not you need them, ask your network administrator.

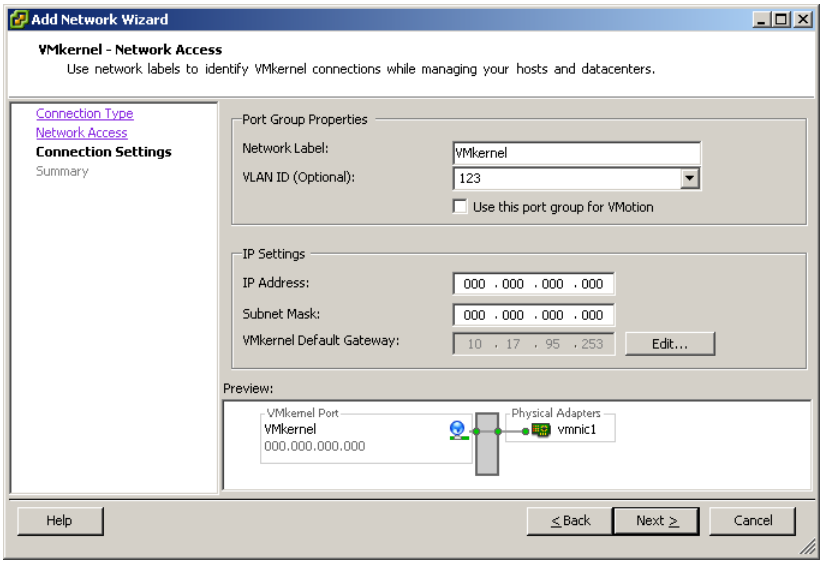

9 In the **IP Settings** group, click **Edit** to set the **VMkernel Default Gateway** for iSCSI.

On the **Routing** tab, the service console and the VMkernel each need their own gateway information.

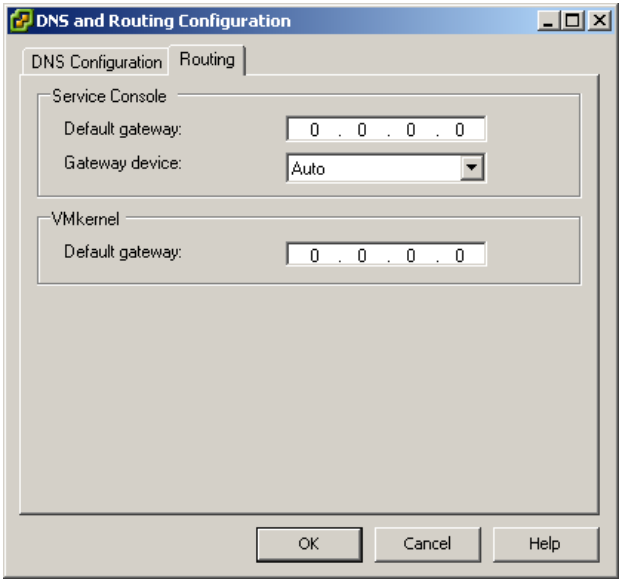

NOTE Set a default gateway for the port that you created. You must use a valid static IP address to configure the VMkernel stack.

10 Click **OK**.

11 Click **Next**.

12 Click **Back** to make any changes.

13 Review your changes on the **Ready to Complete** page and click **Finish**.

After you create a VMkernel port for iSCSI, you must create a service console connection on the same vSwitch as the VMkernel port.

To configure a service console connection for software iSCSI storage

1 Log in to the VI Client and select the server from the inventory panel.

The hardware configuration page for this server appears.

2 Click the **Configuration** tab and click **Networking**.

3 On the right side of the page, click **Properties** for the vSwitch associated with the VMkernel port you have created.

4 On the **Ports** tab, click **Add**.

5 As a connection type, select **Service Console** and click **Next**.

6 In the **Port Group Properties** group, enter a network label that identifies the port group that you are creating.

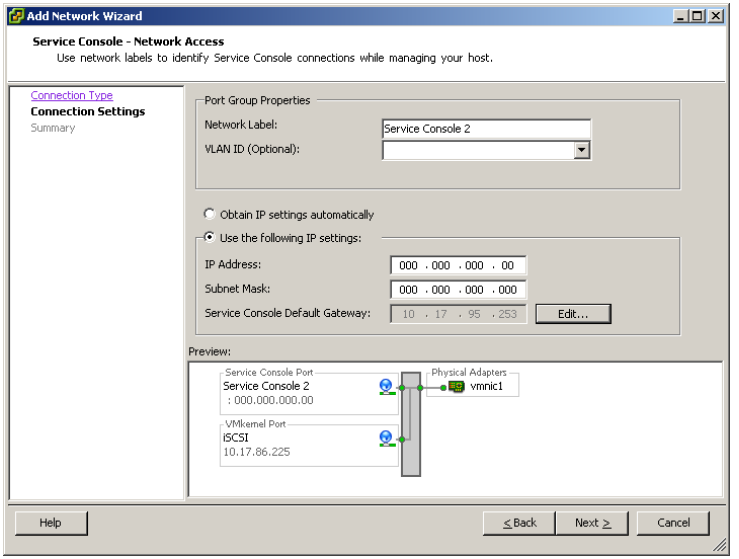

Newer ports and port groups appear at the top of the vSwitch diagram.

7 Enter the **IP Address** and **Subnet Mask,** or select the DHCP option **Obtain IP setting automatically** for the IP address and subnet mask. Note that this must be a different IP than the one chosen for the VMkernel.

8 Click **Edit** to set the **Service Console Default Gateway**.

See "To set the default gateway" on page 37.

9 Click **Next**.

10 After you determine that the vSwitch is configured correctly, click **Finish**.

After you create a VMkernel port and service console connection, you can enable and configure software iSCSI storage. For information on configuring iSCSI adapters and storage, see "iSCSI Storage" on page 110.

Configuring Networking on Blade Servers

Because blade servers might have a limited number of network adapters, you might need to use VLANs to separate traffic for the service console, VMotion, IP storage, and various groups of virtual machines. VMware best practices recommend that the service console and VMotion have their own networks for security reasons. If you dedicate physical adapters to separate vSwitches for this purpose, you might need to relinquish redundant (teamed) connections or stop isolating the various networking clients, or both. VLANs allow you to achieve network segmentation without having to use multiple physical adapters.

For the network blade of a blade server to support an ESX Server 3 port group with VLAN tagged traffic, you must configure the blade to support 802.1Q and configure the port as a tagged port.

The method for configuring a port as a tagged port differs from server to server. The following list describes how to configure a tagged port on three of the most commonly used blade servers:

HP Blade	Set the port's **VLAN Tagging** to **enabled**.
Dell PowerEdge	Set the port to **Tagged**.
IBM eServer Blade Center	Select **Tag** in the port's configuration.

To configure a virtual machine port group with VLAN on a blade server

1 Log in to the VI Client and select the server from the inventory panel.

The hardware configuration page for this server appears.

2 Click the **Configuration** tab, and click **Networking**.

3 On the right side of the page, click **Properties** for vSwitch associated with the service console.

4 On the **Ports** tab, click **Add**.

5 Select **Virtual Machines** for the connection type (default).

6 Click **Next**.

7 In the **Port Group Properties** group, enter a network label that identifies the port group that you are creating.

Use network labels to identify migration-compatible connections common to two or more hosts.

8 In the **VLAN ID** field, enter a number between 1 and 4094.

If you are unsure what to enter, leave this blank or ask your network administrator.

9 Click **Next**.

10 After you determine that the vSwitch is configured correctly, click **Finish**.

To configure a VMkernel port with VLAN on a blade server

1 Log in to the VI Client and select the server from the inventory panel.

The hardware configuration page for this server appears.

2 Click the **Configuration** tab, and click **Networking**.

3 On the right side of the page, click **Properties** for the vSwitch associated with the service console.

4 On the **Ports** tab, click **Add**.

5 Select **VMkernel** and click **Next**.

This option lets you connect the physical network to the VMkernel, which runs services for VMotion and IP storage (NFS or iSCSI).

6 In the **Port Group Properties** group, select or enter a network label and a VLAN ID.

Network Label. A name that identifies the port group that you are creating. This is the label that you specify when configuring a virtual adapter to be attached to this port group, when configuring VMkernel services, such as VMotion and IP storage.

VLAN ID. Identifies the VLAN that the port group's network traffic will use.

7 Select **Use this port group for VMotion** to enable this port group to advertise itself to another ESX Server 3 as the network connection where VMotion traffic should be sent.

You can enable this property for only one VMotion and IP storage port group for each ESX Server 3 host. If this property is not enabled for any port group, migration with VMotion to this host is not possible.

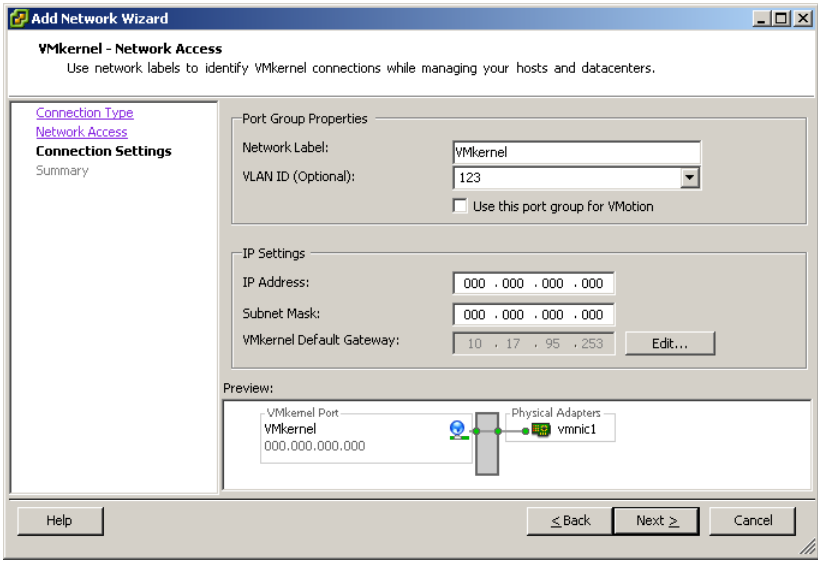

8 In the **IP Settings** group, click **Edit** to set the **VMkernel Default Gateway** for VMkernel services, such as VMotion, NAS, and iSCSI.

NOTE Set a default gateway for the port that you created. VirtualCenter 2 behaves differently than VirtualCenter 1.x. You must use a valid IP address to configure the VMkernel IP stack, not a dummy address.

Under the **DNS Configuration** tab, the name of the host is entered into the name field by default. The DNS server addresses and the domain that were specified during installation are also preselected.

On the **Routing** tab, the service console and the VMkernel each need their own gateway information. A gateway is needed if connectivity to machines not on the same IP subnet as the service console or VMkernel.

Static IP settings is the default.

9 Click **OK**.

10 Click **Next**.

11 Click **Back** to make any changes.

12 Review your changes on the **Ready to Complete** page and click **Finish**.

Troubleshooting

This section guides you through troubleshooting common networking issues.

Troubleshooting Service Console Networking

If certain parts of the service console's networking are misconfigured, you cannot access your ESX Server 3 host with the VI Client. If this happens, you can reconfigure networking by connecting directly to the service console and using the following service console commands:

- `esxcfg-vswif -l`

  Provides a list of the service console's current network interfaces.

  Check that `vswif0` is present and that the current IP address and Netmask are correct.

- `esxcfg-vswitch -l`

  Provides a list of current virtual switch configurations.

  Check that the uplink adapter configured for the service console is connected to the appropriate physical network.

- `exscfg-nics -l`

  Provides a list of current network adapters.

  Check that the uplink adapter configured for the service console is up and that the speed and duplex are both correct.

- `esxcfg-nics -s <speed> <nic>`

  Changes the speed of a network adapter.

- `esxcfg-nics -d <duplex> <nic>`

  Changes the duplex of a network adapter.

- `esxcfg-vswif -i <new ip address> vswifX`

  Changes the service console's IP address.

- `esxcfg-vswif -n <new netmask> vswifX`

  Changes the service console's netmask.

- `esxcfg-vswitch -U <old vmnic> <service console vswitch>`

  Removes the uplink for the service console

- `esxcfg-vswitch -L <new vmnic> <service console vswitch>`

  Changes the uplink for the service console.

If you encounter long waits when using `esxcfg-*` commands, the DNS might be misconfigured. The `esxcfg-*` commands require that DNS be configured so that localhost name resolution works properly. This requires that the `/etc/hosts` file contain an entry for the configured IP address and the 127.0.0.1 localhost address.

Troubleshooting Network Adapter Configuration

Adding a new network adapter, in certain cases, can cause loss of service console connectivity and manageability by using the VI Client because of network adapters getting renamed.

If this happens, you must use the service console to rename the affected network adapters.

To rename network adapters by using the service console

1. Log in directly to your ESX Server 3 host's console.
2. Use the `esxcfg-nics -l` command to see which names were assigned to your network adapters.
3. Use the `esxcfg-vswitch -l` command to see which vSwitches, if any, are now associated with device names no longer shown by esxcfg-nics.
4. Use the `esxcfg-vswitch -U <old vmnic name> <vswitch>` command to remove any network adapters that were renamed.
5. Use the `esxcfg-vswitch -L <new vmnic name> <vswitch>` command to re-add the network adapters, giving them the correct names.

Troubleshooting Physical Switch Configuration

In some cases, you might lose vSwitch connectivity when a failover or failback event occurs. This causes the MAC addresses that virtual machines associated with that vSwitch use to appear on a different switch port than they previously did.

To avoid this problem, put your physical switch in portfast or portfast trunk mode.

Troubleshooting Port Group Configuration

Changing the name of a port group when virtual machines are already connected to that port group causes invalid network configuration for the virtual machines configured to connect to that port group.

The connection from virtual network adapters to port groups is made by name, and the name is what is stored in the virtual machine configuration. Changing the name of a port group does not cause a mass reconfiguration of all the virtual machines connected to that port group. Virtual machines that are already powered on continue to function until they are powered off because their connections to the network are already established.

Avoid renaming networks after they are in use. After you rename a port group, you must reconfigure each associated virtual machine by using the service console to reflect the new port group name.

Storage

Introduction to Storage 5

The Storage section contains overview information about available storage options for ESX Server 3 and explains how to configure you ESX Server 3 system so that it can use and manage different types of storage.

For information on specific activities that a storage administrator might need to perform on storage arrays, see the *Fibre Channel SAN Configuration Guide* and the *iSCSI SAN Configuration Guide.*

This chapter discusses the following topics:

- "Storage Overview" on page 88
- "Types of Physical Storage" on page 88
- "Supported Storage Adapters" on page 91
- "Datastores" on page 91
- "Comparing Types of Storage" on page 97
- "Viewing Storage Information in the VMware Infrastructure Client" on page 97
- "Configuring and Managing Storage" on page 101

Storage Overview

An ESX Server 3 virtual machine uses a virtual hard disk to store its operating system, program files, and other data associated with its activities. A virtual disk is a large physical file, or a set of files, that can be copied, moved, archived, and backed up as easily as any other file. To store virtual disk files and manipulate the files, ESX Server 3 requires dedicated storage space.

ESX Server 3 uses storage space on a variety of physical storage devices, including your host's internal and external storage devices, or networked storage devices. The storage device is a physical disk or disk array dedicated to the specific tasks of storing and protecting data.

ESX Server 3 can discover storage devices that it has access to and format them as datastores. The datastore is a special logical container, analogous to a file system on a logical volume, where ESX Server 3 places virtual disk files and other files that encapsulate essential components of a virtual machine. Deployed on different devices, the datastores hide specifics of each storage product and provide a uniform model for storing virtual machine files.

Using the VI Client, you can set up datastores in advance on any storage device that your ESX Server 3 discovers.

To learn how to access and configure your storage devices, as well as how to create and manage datastores, see the following chapters:

- Chapter 6, "Configuring Storage," on page 103
- Chapter 7, "Managing Storage," on page 133

After you create the datastores, you can use them to store virtual machine files. For information on creating virtual machines, see *Basic System Administration*.

Types of Physical Storage

The ESX Server 3 storage management process starts with a storage space that your storage administrator preallocates on different storage devices.

ESX Server 3 supports the following types of storage devices:

- **Local** – Stores virtual machine files on internal or external storage devices or arrays attached to your ESX Server 3 host through a direct connection.
- **Networked** – Stores virtual machine files on external shared storage devices or arrays located outside of your ESX Server 3 host. The host communicates with the networked devices through a high-speed network.

Local Storage

Local storage devices can be internal hard disks located inside your ESX Server 3 host, or external storage systems, located outside and connected to the host directly.

Local storage devices do not require a storage network to communicate with your ESX Server 3. All you need is a cable connected to the storage device and, when required, a compatible HBA in your ESX Server 3 host.

Generally, you can connect multiple ESX Server 3 hosts to a single local storage system. The actual number of hosts you connect varies depending on the type of storage device and topology you use.

Many local storage systems support redundant connection paths to ensure fault tolerance. See "Managing Multiple Paths" on page 137.

When multiple ESX Server 3 hosts connect to the local storage unit, they access storage LUNs in the unshared mode. The unshared mode does not permit several ESX Server 3 hosts to access the same VMFS datastore concurrently. However, a few SAS storage systems offer shared access to multiple ESX Server 3i hosts. This type of access permits multiple ESX Server 3i hosts to access the same VMFS datastore on a LUN. See "Sharing a VMFS Volume Across ESX Server 3 Systems" on page 94.

ESX Server 3 supports a variety of internal or external local storage devices, including SCSI, IDE, SATA, and SAS storage systems. No matter which type of storage you use, ESX Server 3 hides a physical storage layer from virtual machines.

When you set up your local storage, keep in mind the following:

- You cannot use IDE/ATA drives to store virtual machines.
- Use local SATA storage, internal and external, in unshared mode only. SATA storage does not support sharing the same LUNs and, therefore, the same VMFS datastore across multiple ESX Server 3 hosts.

  When you use SATA storage, ensure that your SATA drives are connected through supported dual SATA/SAS controllers.
- Some SAS storage systems can offer shared access to the same LUNs (and, therefore, the same VMFS datastores) to multiple ESX Server 3 hosts.

For information on supported local storage devices, see the *I/O Compatibility Guide* at www.vmware.com/support/pubs/vi_pubs.html.

Networked Storage

Networked storage devices are external storage devices, or arrays, that your ESX Server 3 host uses to store virtual machine files remotely. The ESX Server 3 host accesses these devices over a high-speed storage network.

ESX Server 3 supports the following networked storage technologies:

- **Fibre Channel (FC)** – Stores virtual machine files remotely on an FC storage area network (SAN). FC SAN is a specialized high-speed network that connects your ESX Server 3 hosts to high-performance storage devices. The network uses Fibre Channel protocol to transport SCSI traffic from virtual machines to the FC SAN devices.

  To connect to the FC SAN, your ESX Server 3 host should be equipped with Fibre Channel host bus adapters (HBAs) and, unless you use Fibre Channel direct connect storage, with Fibre Channel switches that help route storage traffic.

- **Internet SCSI (iSCSI)** – Stores virtual machine files on remote iSCSI storage devices. iSCSI packages SCSI storage traffic into the TCP/IP protocol so that it can travel through standard TCP/IP networks instead of the specialized FC network. With iSCSI connection, your ESX Server 3 host serves as the *initiator* that communicates with a *target*, located in remote iSCSI storage systems.

  ESX Server 3 offers the following types of iSCSI connection:

  - **Hardware Initiated iSCSI** – Your ESX Server 3 host connects to storage through a third-party iSCSI HBA.
  - **Software Initiated iSCSI** – Your ESX Server 3 uses a software-based iSCSI initiator in the VMkernel to connect to storage. With this type of iSCSI connection, your host needs only a standard network adapter for network connectivity.

- **Network-Attached Storage (NAS)** – Stores virtual machine files on remote file servers accessed over standard TCP/IP network. The NFS client built into ESX Server 3 uses the Network File System (NFS) protocol version 3 to communicate with the NAS/NFS servers. For network connectivity, the ESX Server 3 host requires a standard network adapter.

See the *Storage/SAN Compatibility Guide* at www.vmware.com/pdf/vi3_san_guide.pdf.

Supported Storage Adapters

Depending on the type of storage available to you, your ESX Server 3 system might need adapters that provide connectivity to a specific storage device or network. ESX Server 3 supports different classes of adapters, including SCSI, iSCSI, RAID, Fibre Channel, and Ethernet. ESX Server 3 accesses the adapters directly through device drivers in the VMkernel.

For details on the types of adapters that ESX Server 3 supports, see the *I/O Compatibility Guide* at www.vmware.com/support/pubs/vi_pubs.html.

Datastores

You use the VI Client to access different types of storage devices that your ESX Server 3 host discovers and to deploy datastores on them. Datastores are special logical containers, analogous to file systems, that hide specifics of each storage device and provide a uniform model for storing virtual machine files.

Datastores can also be used for storing ISO images, virtual machine templates, and floppy images. See *Basic System Administration* at www.vmware.com/support/pubs/.

Depending on the type of storage you use, ESX Server 3 datastores can have the following file system formats:

- **Virtual Machine File System (VMFS)** – Special high-performance file system optimized for storing ESX Server 3 virtual machines. ESX Server 3 can deploy VMFS on any SCSI-based local or networked storage device, including Fibre Channel and iSCSI SAN equipment.

  As an alternative to using the VMFS datastore, your virtual machine can have direct access to raw devices and use a mapping file (RDM) as a proxy. For more information on RDMs, see "Raw Device Mapping" on page 145.

- **Network File System (NFS)** – File system on a NAS storage device. ESX Server 3 supports NFS version 3 over TCP/IP. ESX Server 3 can access a designated NFS volume located on an NFS server. ESX Server 3 mounts the NFS volume and uses it for its storage needs.

If you use service console to access your ESX Server 3 host, you can see the VMFS and NFS datastores as separate subdirectories in the `/vmfs/volumes` directory. For information on using the service console commands and utilities, see "Using vmkfstools" on page 283.

VMFS Datastores

When your ESX Server 3 host accesses SCSI-based storage devices such as SCSI, iSCSI, or FC SAN, the storage space is presented to your ESX Server 3 as a LUN. A *LUN* is a logical volume that represents storage space on a single physical disk or on a number of disks aggregated in a disk array. A single LUN can be created from the entire space on the storage disk or array, or from a part of the space, called a *partition*. The LUN that uses disk space on more than one physical disk or partition still presents itself as a single logical volume to your ESX Server 3.

ESX Server 3 can format LUNs as VMFS datastores. VMFS datastores primarily serve as repositories for virtual machines. You can store multiple virtual machines on the same VMFS volume. Each virtual machine, encapsulated in a set of files, occupies a separate single directory. For the operating system inside the virtual machine, VMFS preserves the internal file system semantics, which ensures correct application behavior and data integrity for applications running in virtual machines.

In addition, you can use the VMFS datastores to store other files, such as virtual machine templates and ISO images.

VMFS supports the following file and block sizes enabling your virtual machines to run even the most data intensive applications, including databases, ERP, and CRM in virtual machines:

- Maximum virtual disk size: 2TB with 8MB block size
- Maximum file size: 2TB with 8MB block size
- Block size: 1MB (default), 2MB, 4MB, and 8MB

Creating and Increasing VMFS Datastores

You use the VI Client to set up a VMFS datastore in advance on any SCSI-based storage device that your ESX Server 3 discovers. ESX Server 3 lets you have up to 256 VMFS datastores per system with the minimum volume size 1.2GB.

NOTE Always have only one VMFS datastore for each LUN.

For information on creating VMFS datastores on the SCSI-based storage devices, see the following sections:

- "Adding Local Storage" on page 104
- "Adding Fibre Channel Storage" on page 108
- "Adding iSCSI Storage Accessible Through Hardware Initiators" on page 119
- "Adding iSCSI Storage Accessible Through Hardware Initiators" on page 119

After you create the VMFS datastore, you can edit its properties. See "Editing VMFS Datastores" on page 134.

If your VMFS datastore requires more space, you can dynamically increase the VMFS volume, up to 64TB, by adding an extent. An *extent* is a LUN on a physical storage device that can be dynamically added to any existing VMFS datastore. The datastore can stretch over multiple extents, yet appear as a single volume.

NOTE You cannot reformat a VMFS volume that a remote ESX Server 3 host is using. If you attempt to do so, you receive a warning to this effect that specifies the name of the volume in use and the MAC address of a host NIC that is using it. This warning also appears in the VMkernel and vmkwarning log files.

Considerations when Creating VMFS Datastores

Plan how to set up storage for your ESX Server 3 systems before you format storage devices with a VMFS datastore.

You might want fewer, larger VMFS volumes for the following reasons:

- More flexibility to create virtual machines without asking the storage administrator for more space.
- More flexibility for resizing virtual disks, doing snapshots, and so on.
- Fewer VMFS datastores to manage.

You might want more, smaller VMFS volumes for the following reasons:

- Less wasted storage space.
- Different applications might need different RAID characteristics.
- More flexibility, as the multipathing policy and disk shares are set per LUN.
- Use of Microsoft Cluster Service requires that each cluster disk resource is in its own LUN.
- Better performance.

You might decide to configure some of your servers to use fewer, larger VMFS volumes and other servers to use more, smaller VMFS volumes.

Sharing a VMFS Volume Across ESX Server 3 Systems

As a cluster file system, VMFS lets multiple ESX Server 3 hosts access the same VMFS datastore concurrently. You can connect up to 32 hosts to a single VMFS volume.

Figure 5-1. Sharing a VMFS Volume Across ESX Server 3 Hosts

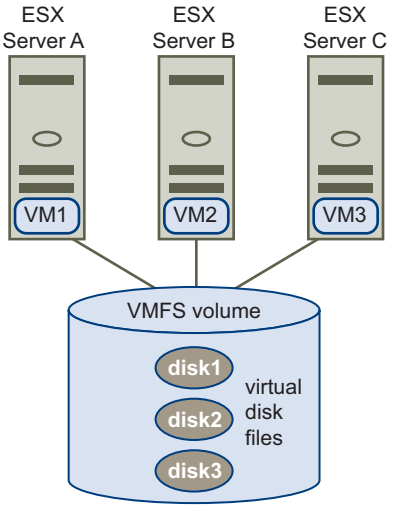

To ensure that the same virtual machine is not accessed by multiple servers at the same time, VMFS provides on-disk locking.

Sharing the same VMFS volume across multiple ESX Server 3 hosts gives you the following advantages:

- You can use VMware Distributed Resource Scheduling and VMware High Availability.

  You can distribute virtual machines across different physical servers. That means you run a mix of virtual machines on each server so that not all experience high demand in the same area at the same time.

  If a server fails, you can restart virtual machines on another physical server. In case of a failure, the on-disk lock for each virtual machine is released.

  For more information on VMware DRS and VMware HA, see the *Resource Management Guide* at www.vmware.com/support/pubs/.

- You can use VMotion to perform live migration of running virtual machines from one physical server to another.

  For more information on VMotion, see *Basic System Administration* at www.vmware.com/support/pubs/.

- You can use VMware Consolidated Backup, which lets a proxy server, called VCB proxy, back up a snapshot of a virtual machine while the virtual machine is powered on and is reading and writing to its storage.

  For more information on Consolidated Backup, see the *Virtual Machine Backup Guide* at www.vmware.com/support/pubs/.

NFS Datastore

ESX Server 3 can access a designated NFS volume located on a NAS server, mount this volume, and use it for its storage needs. You can use NFS volumes to store and boot virtual machines in the same way you use VMFS datastores.

ESX Server 3 supports the following shared storage capabilities on NFS volumes:

- Use VMotion.
- Use VMware DRS and VMware HA.
- Mount ISO images, which are presented as CD-ROMs to virtual machines.
- Create virtual machine snapshots. See *Basic System Administration* at www.vmware.com/support/pubs/.

How Virtual Machines Access Storage

When a virtual machine communicates with its virtual disk stored on a datastore, it issues SCSI commands. Because datastores can exist on various types of physical storage, these commands are encapsulated into other forms, depending on the protocol that the ESX Server 3 system uses to connect to a storage device. ESX Server 3 supports Fibre Channel (FC), Internet SCSI (iSCSI), and NFS protocols.Regardless of the type of storage device your ESX Server 3 uses, the virtual disk always appears to the virtual machine as a mounted SCSI device. The virtual disk hides a physical storage layer from the virtual machine's operating system. This allows you to run even operating systems that are not certified for specific storage equipment, such as SAN, inside the virtual machine.

Figure 5-2 depicts five virtual machines using different types of storage to illustrate the differences between each type.

Figure 5-2. Virtual Machines Accessing Different Types of Storage

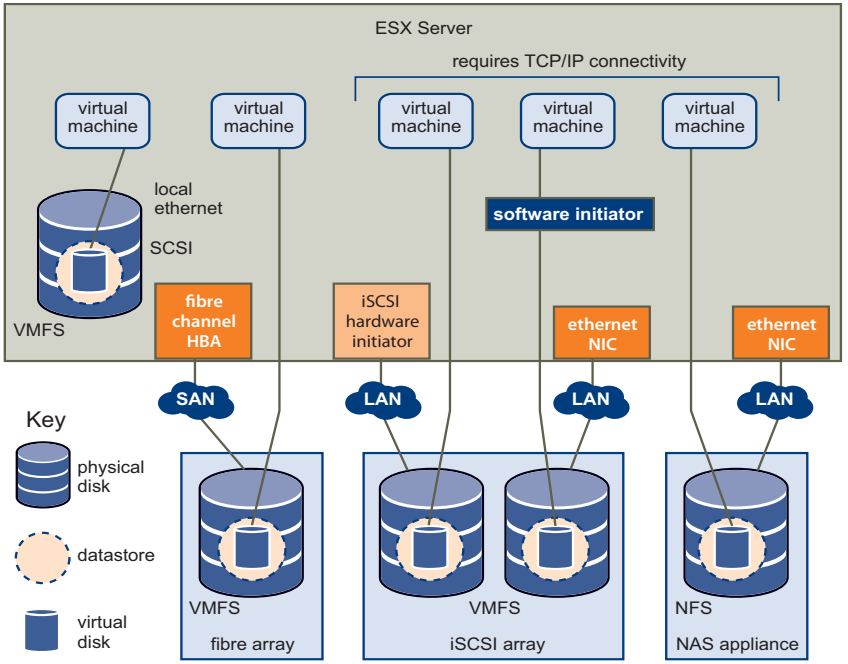

NOTE This diagram is for conceptual purposes only. It is not a recommended configuration.

Comparing Types of Storage

Table 5-1 compares networked storage technologies ESX Server 3 supports.

Table 5-1. Networked Storage that ESX Server 3 Supports

Technology	Protocols	Transfers	Interface
Fibre Channel	FC/SCSI	Block access of data/LUN	FC HBA
iSCSI	IP/SCSI	Block access of data/LUN	■ iSCSI HBA (hardware-initiated iSCSI) ■ NIC (software-initiated iSCSI)
NAS	IP/NFS	File (no direct LUN access)	NIC

Table 5-2 compares the ESX Server 3 features that different types of storage support.

Table 5-2. ESX Server 3 Features Supported by Storage

Storage Type	Boot VM	VMotion	Datastore	RDM	VM Cluster	VMware HA and DRS	VCB
SCSI	Yes	No	VMFS	No	No	No	Yes
Fibre Channel	Yes	Yes	VMFS	Yes	Yes	Yes	Yes
iSCSI	Yes	Yes	VMFS	Yes	No	Yes	Yes
NAS over NFS	Yes	Yes	NFS	No	No	Yes	Yes

Viewing Storage Information in the VMware Infrastructure Client

The VI Client displays detailed information on available datastores, storage devices that the datastores use, and configured adapters. For more information, see these sections:

- "Displaying Datastores" on page 98
- "Viewing Storage Adapters" on page 99
- "Understanding Storage Device Naming in the Display" on page 100

Displaying Datastores

Datastores are added to the VI Client in the following ways:

- Discovered when a host is added to the inventory. When you add a host to the inventory, the VI Client displays any datastores available to the host.
- Created on an available storage device. You can use the **Add Storage** option to create and configure a new datastore. See "Configuring Storage" on page 103.

You can view a list of available datastores and analyze their properties.

To display datastores, on the host **Configuration** tab, click **Storage**.

For each datastore, the Storage section shows summary information, including:

- Target storage device where the datastore is located. See "Understanding Storage Device Naming in the Display" on page 100.
- Type of file system the datastore uses. See "Datastores" on page 91.
- Total capacity, including the used and available space.

To view additional details about the specific datastore, select the datastore from the list. The **Details** pane shows the following information:

- Location of the datastore.
- Individual extents that the datastore spans and their capacity (VMFS datastores).
- Paths used to access the storage device (VMFS datastores).

In Figure 5-3, the symm-07 datastore is selected from the list of available datastores. The **Details** pane provides information about the selected datastore.

Figure 5-3. Datastore Information

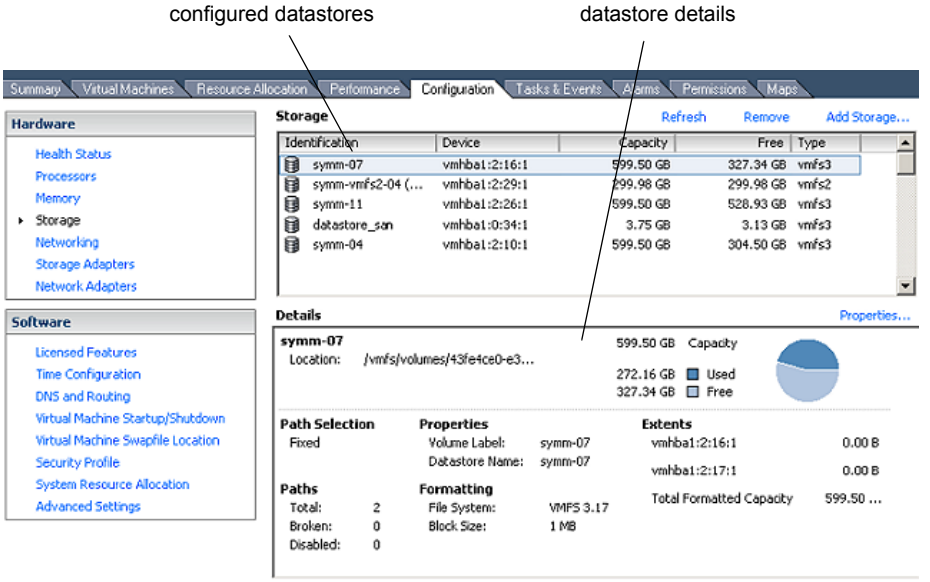

You can refresh and remove any of the existing datastores, and change the properties of a VMFS datastore. When you edit or reconfigure a VMFS datastore, you can change its label, add extents, upgrade it, or modify paths for storage devices. See "Managing Storage" on page 133.

Viewing Storage Adapters

The VI Client displays any storage adapters available to your system.

To display storage adapters, on the host **Configuration** tab, click **Storage Adapters**.

You can view the following information about the storage adapters:

- Existing storage adapters.
- Type of storage adapter, such as Fibre Channel SCSI or iSCSI.
- Details for each adapter, such as the storage device it connects to and target ID.

To view configuration properties for a specific adapter, select the adapter from the **Storage Adapters** list.

In Figure 5-4, the iSCSI hardware vmhba0 adapter is selected. The **Details** pane provides information about the number of LUNs the adapter connects to and the paths it uses.

To change the path's configuration, select this path from the list, right-click the path, and click **Manage Paths** to open the Manage Paths dialog box. See "Managing Multiple Paths" on page 137.

Figure 5-4. Storage Adapter Information

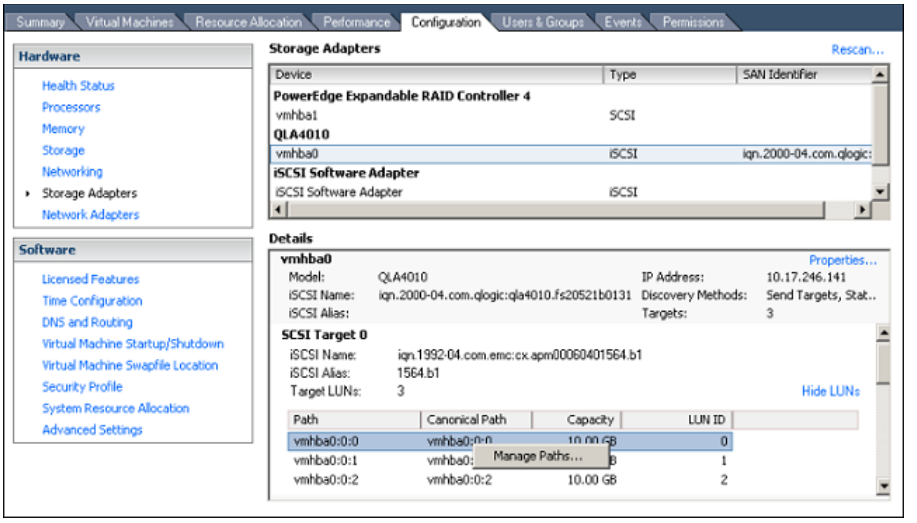

Understanding Storage Device Naming in the Display

In the VI Client, the name of a storage device is displayed as a sequence of three or four numbers, separated by colons, such as `vmhba1:1:3:1`. The name has the following meaning:

```
<HBA>:<SCSI target>:<SCSI LUN>:<disk partition>
```

The abbreviation `vmhba` refers to different physical HBAs on the ESX Server 3 system. It can also refer to the virtual iSCSI initiator that ESX Server 3 implements by using the VMkernel network stack. The fourth number indicates a partition on a disk that a VMFS datastore occupies.

The `vmhba1:1:3:1` example refers to the first partition on SCSI LUN3, SCSI target 1, which is accessed through HBA 1.

Although the third and the fourth numbers never change, the first two numbers can change. For example, after rebooting the ESX Server 3 system, `vmhba1:1:3:1` can change to `vmhba3:2:3:1`, however, the name still refers to the same physical device. The first and the second numbers can change for the following reasons:

- The first number, the HBA, changes when an outage on the Fibre Channel or iSCSI network occurs. In this case, the ESX Server 3 system must use a different HBA to access the storage device.
- The second number, the SCSI target, changes if any modifications occur in the mappings of the Fibre Channel or iSCSI targets visible to the ESX Server 3 host.

Configuring and Managing Storage

The Configuring Storage and Managing Storage chapters of this guide discuss most of the concepts and outline tasks you need to perform when working with storage.

For detailed information on configuring SANs, see the *Fibre Channel SAN Configuration Guide* or *iSCSI SAN Configuration Guide*.

For more information about specific storage configuration tasks, see the following:

- Local storage configuration:

  "To create a datastore on a local SCSI disk" on page 105

- Fibre Channel SAN storage configuration:

  "To create a datastore on a Fibre Channel device" on page 108

- Hardware-initiated iSCSI storage configuration:
  - "To view the hardware iSCSI initiator properties" on page 113
  - "To set up the iSCSI name, alias, and IP address for the hardware initiator" on page 115
  - "To set up target discovery addresses using dynamic discovery" on page 116
  - "To set up CHAP parameters for the hardware initiator" on page 118
  - "To create a datastore on a hardware iSCSI device" on page 119
- Software-initiated iSCSI storage configuration:
  - "To view the software iSCSI initiator properties" on page 121
  - "To enable the software iSCSI initiator" on page 124
  - "To set up target discovery addresses for the software initiator" on page 124

- "To set up CHAP parameters for the software initiator" on page 124
- "To create a datastore on an iSCSI device accessed through software initiators" on page 125

- NAS storage configuration:

  "To mount an NFS volume" on page 129

- Storage management:
  - "To upgrade the VMFS-2 to VMFS-3" on page 135
  - "To edit the name of the datastore" on page 136
  - "To add one or more extents to the datastore" on page 136
  - "To remove a datastore" on page 134
- Path managing:
  - "To set the multipathing policy" on page 143
  - "To set the preferred path" on page 144
  - "To disable a path" on page 144

6 Configuring Storage

This chapter contains information about configuring local SCSI storage devices, Fibre Channel SAN storage, iSCSI storage, and NAS storage.

NOTE For additional information about configuring SANs, see the *Fibre Channel SAN Configuration Guide* and *iSCSI SAN Configuration Guide*.

This chapter discusses the following topics:

- "Local Storage" on page 104
- "Fibre Channel Storage" on page 107
- "iSCSI Storage" on page 110
- "Performing a Rescan" on page 126
- "Network Attached Storage" on page 127
- "Creating a Diagnostic Partition" on page 130

Local Storage

Local storage uses a SCSI-based device such as your ESX Server 3 host's hard disk or any external dedicated storage system connected directly to your ESX Server 3 host. Figure 6-1 depicts a virtual machine using local SCSI storage.

Figure 6-1. Local Storage

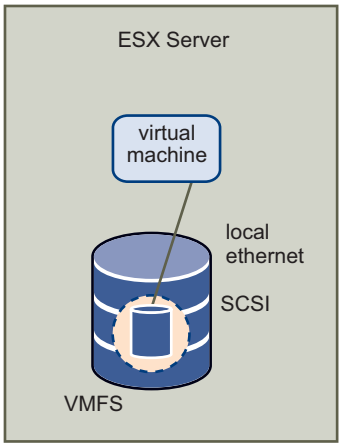

In this example of local storage topology, the ESX Server 3 host uses a single connection to plug into a disk. On that disk, you can create a VMFS datastore, which you use to store virtual machine disk files.

Although this storage configuration is possible, it is not a recommended topology. Using single connections between storage arrays and ESX Server 3 hosts creates *single points of failure (SPOF)* that can cause interruptions when a connection becomes unreliable or fails. To ensure fault tolerance, some DAS systems support redundant connection paths. See "Managing Multiple Paths" on page 137.

Adding Local Storage

As soon as you load storage adapter drivers, ESX Server 3 detects available SCSI storage devices. Before you create a new datastore on a SCSI device, you might need to perform a rescan. See "Performing a Rescan" on page 126.

When you create a datastore on a SCSI storage device, the Add Storage wizard guides you through the configuration steps.

To create a datastore on a local SCSI disk

1 Log in to the VI Client and select the server from the inventory panel.

2 Click the **Configuration** tab and click **Storage** in the **Hardware** panel.

3 Click **Add Storage**.

4 Select the **Disk/LUN** storage type and click **Next**.

5 Select the SCSI device to use for your datastore and click **Next**.

The Current Disk Layout page opens. If the disk you are formatting is blank, the Current Disk Layout page automatically presents the entire disk space to you for storage configuration.

6 If the disk is not blank, review the current disk layout in the top panel of the Current Disk Layout page and select a configuration option from the bottom panel:

- **Use the entire device** — Select this option to dedicate the entire disk or LUN to a single VMFS datastore. VMware recommends that you select this option.

WARNING If you select this option, any file systems or data previously stored on this device will be destroyed.

- **Use free space** — Select this option to deploy a VMFS datastore in the remaining free space of the disk.

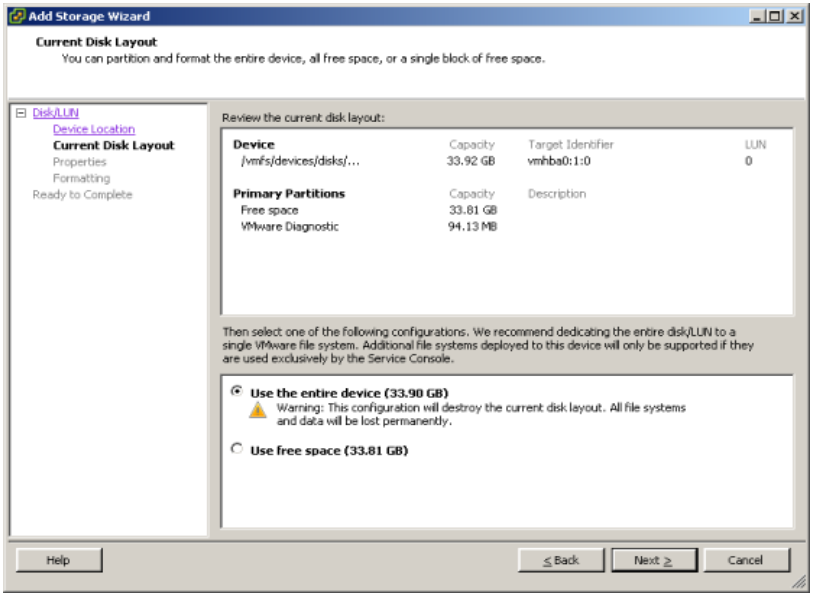

7 Click **Next**.

8 In the Disk/LUN–Properties page, enter a datastore name and click **Next**.

The Disk/LUN–Formatting page appears.

9 If needed, adjust the file system and capacity values.

By default, the entire free space available on the storage device is offered to you.

10 Click **Next**.

11 In the Ready to Complete page, review the datastore configuration information and click **Finish**.

This process creates a datastore on the local SCSI disk on your ESX Server 3 host.

Fibre Channel Storage

ESX Server 3 supports Fibre Channel adapters, which allow an ESX Server 3 system to be connected to a SAN and see the disk arrays on the SAN.

Figure 6-2 depicts virtual machines using Fibre Channel storage.

Figure 6-2. Fibre Channel Storage

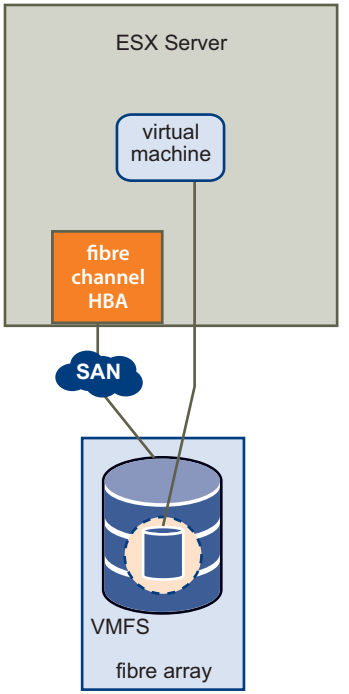

In this configuration, an ESX Server 3 system connects to SAN fabric, which consists of Fibre Channel switches and storage arrays, using a Fibre Channel adapter. LUNs from a storage array become available to your ESX Server 3 system. You can access the LUNs and create a datastore that you use for your ESX Server 3 storage needs. The datastore uses the VMFS format.

For additional information:

- See "Adding Fibre Channel Storage" on page 108.
- About configuring SANs, see the *Fibre Channel SAN Configuration Guide*.

- About supported SAN storage devices for ESX Server 3, see the *Storage/SAN Compatibility Guide*.
- About multipathing for Fibre Channel HBAs and how to manage paths, see "Managing Multiple Paths" on page 137.

Adding Fibre Channel Storage

Before you create a new datastore on a Fibre Channel device, rescan a Fibre Channel adapter to discover any newly added LUNs. See "Performing a Rescan" on page 126.

When you create a datastore on a Fibre Channel storage device, the Add Storage wizard guides you through the configuration.

To create a datastore on a Fibre Channel device

1. Log in to the VI Client, and select a server from the inventory panel.
2. Click the **Configuration** tab and click **Storage** in the **Hardware** panel.
3. Click **Add Storage**.
4. Select the **Disk/LUN** storage type and click **Next**.
5. Select the Fibre Channel device to use for your datastore, and click **Next**.

   The Current Disk Layout page opens. If the disk you are formatting is blank, the Current Disk Layout page automatically presents the entire disk space to you for storage configuration.
6. If the disk is not blank, review the current disk layout in the top panel of the Current Disk Layout page and select a configuration option from the bottom panel:
   - **Use the entire device** — Select this option to dedicate the entire disk or LUN to a single VMFS datastore. VMware recommends that you select this option.

WARNING If you select this option, any file systems or data previously stored on this device will be destroyed.

- **Use free space** — Select this option to deploy a VMFS datastore in the remaining free space of the disk.

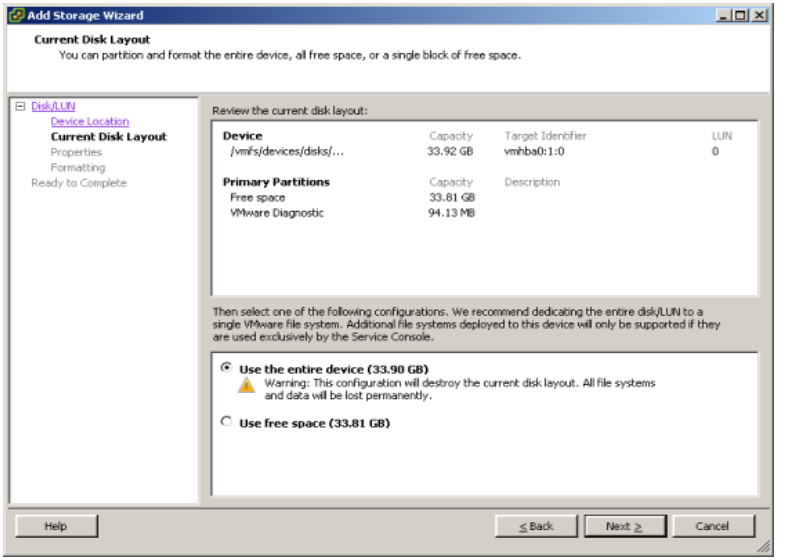

7 Click **Next**.

8 In the Disk/LUN–Properties page, enter a datastore name and click **Next**.

The Disk/LUN–Formatting page appears.

9 If needed, adjust the file system values and capacity you use for the datastore.

By default, the entire free space available on the storage device is offered to you.

10 Click **Next**.

11 In the Ready to Complete page, review the datastore configuration information and click **Finish**.

This process creates the datastore on a Fibre Channel disk for the ESX Server 3 host.

12 Click **Refresh**.

For advanced configuration, such as using multipathing, masking, and zoning, see the *Fibre Channel SAN Configuration Guide*.

iSCSI Storage

ESX Server 3 supports iSCSI technology that allows your ESX Server 3 system to use an IP network while accessing remote storage. With iSCSI, SCSI storage commands that your virtual machine issues to its virtual disk are converted into TCP/IP protocol packets and transmitted to a remote device, or target, that stores the virtual disk. From the point of view of the virtual machine, the device appears as a locally attached SCSI drive.

iSCSI Initiators

To access remote targets, your ESX Server 3 host uses iSCSI initiators. Initiators transport SCSI requests and responses between the ESX Server 3 system and the target storage device on the IP network.

ESX Server 3 supports hardware-based and software-based iSCSI initiators:

- **Hardware iSCSI initiator** – A third-party host bus adapter (HBA) with iSCSI over TCP/IP capability. This specialized iSCSI adapter is responsible for all iSCSI processing and management.
- **Software iSCSI initiator** – A code built into VMkernel that lets your ESX Server 3 system connect to the iSCSI storage device through standard network adapters. The software initiator handles the iSCSI processing while communicating with the network adapter through the network stack. With the software initiator, you can use the iSCSI technology without purchasing specialized hardware.

Figure 6-3 depicts two virtual machines that use different types of iSCSI initiators.

Figure 6-3. iSCSI Storage

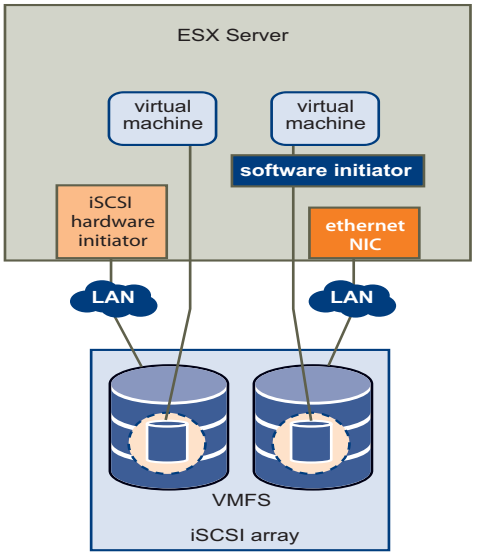

In the first example of iSCSI storage configuration, the ESX Server 3 system uses the hardware iSCSI adapter. This specialized iSCSI adapter sends iSCSI packets to a disk over a LAN.

In the second example, the ESX Server 3 system is configured with the software iSCSI initiator. Using the software initiator, the ESX Server 3 system connects to a LAN through an existing NIC card.

Naming Requirements

Because SANs can become large and complex, all iSCSI initiators and targets that use the network have unique and permanent iSCSI names and are assigned addresses for access. The iSCSI name provides a correct identification of a particular iSCSI device, an initiator or a target, regardless of its physical location.

When you configure your iSCSI initiators, make sure they have properly formatted names. The initiators can use one of the following formats:

- **IQN (iSCSI qualified name)** – Can be up to 255 characters long and has the following format:

  `iqn.<year-mo>.<reversed_domain_name>:<unique_name>`

  where `<year-mo>` represents the year and month your domain name was registered, `<reversed_domain_name>` is the official domain name, reversed, and `<unique_name>` is any name you want to use, for example, the name of your server.

  An example might be `iqn.1998-01.com.mycompany:myserver`.

- **EUI (extended unique identifier)** – Represents the `eui.` prefix followed by the 16-character name. The name includes 24 bits for company name assigned by the IEEE and 40 bits for a unique ID such as a serial number.

  For example, `eui.0123456789ABCDEF`.

Discovery Methods

To determine which storage resource on the network is available for access, the ESX Server 3 system uses these discovery methods:

- **Dynamic discovery** – Also known as Send Targets discovery. Each time the initiator contacts a specified iSCSI server, it sends the *Send Targets* request to the server. The server responds by providing a list of available targets to the initiator.
- **Static Discovery** – The initiator does not need to perform any discovery. The initiator in advance knows all targets it will be contacting and uses their IP addresses and domain names to communicate with them.

  The static discovery method is available only when the iSCSI storage is accessed through hardware initiators.

iSCSI Security

Because iSCSI technology uses the IP networks to connect to remote targets, it is necessary to ensure security of the connection. The IP protocol itself doesn't protect the data it transports, and it doesn't have the capability to verify the legitimacy of initiators that access targets on the network. You need to take specific measures to guarantee security across IP networks.

ESX Server 3 supports the Challenge Handshake Authentication Protocol (CHAP) that your iSCSI initiators can use for authentication purposes. After your initiator establishes the initial connection with the target, CHAP verifies the identity of the initiator and checks a CHAP secret that your initiator and the target share. This can be repeated periodically during the iSCSI session.

When you configure iSCSI initiators for your ESX Server 3 system, check whether the iSCSI storage supports CHAP and if it does, make sure to enable it for your initiators. See "Securing iSCSI Storage" on page 204.

Configuring Hardware iSCSI Initiators and Storage

When your ESX Server 3 communicates with the iSCSI storage through hardware initiators, it uses a specialized third-party adapter that can access iSCSI storage over TCP/IP. This iSCSI adapter handles all iSCSI processing and management for your ESX Server 3 system.

Install and configure the hardware iSCSI adapter before you set up the datastore that resides on an iSCSI storage device.

Installing and Viewing Hardware iSCSI Initiators

For information on which adapters are supported, see the *I/O Compatibility Guide* on the VMware Web site at www.vmware.com.

Before you begin configuring the hardware iSCSI initiator, make sure that the iSCSI HBA is successfully installed and appears on the list of adapters available for configuration. If the initiator is installed, you can view its properties.

To view the hardware iSCSI initiator properties

1 Log in to the VI Client, and select a server from the inventory panel.

2 Click the **Configuration** tab and click **Storage Adapters** in the **Hardware** panel.

The hardware iSCSI initiator appears in the list of storage adapters.

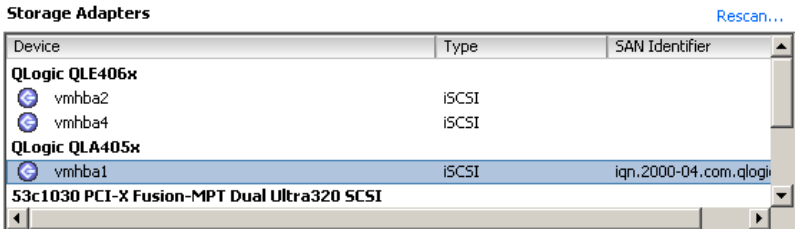

3 Select the initiator to configure.

The details for the initiator appear, including the model, IP address, iSCSI name, discovery methods, iSCSI alias, and any discovered targets.

4 Click **Properties**.

The iSCSI Initiator Properties dialog box opens. The **General** tab displays additional characteristics of the initiator.

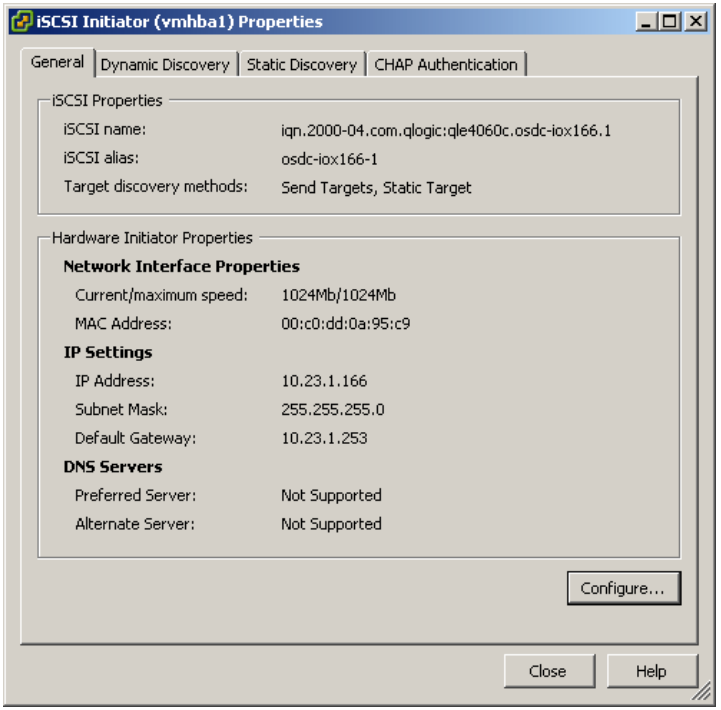

You can now configure your hardware initiator or change its default characteristics.

Configuring Hardware iSCSI Initiators

While you configure the hardware iSCSI initiator, set up your initiator's iSCSI name, IP address, and discovery addresses. In addition, VMware recommends that you set up CHAP parameters.

After you configure your hardware iSCSI initiator, perform a rescan, so that all LUNs that the initiator can access appear on the list of storage devices. See "Performing a Rescan" on page 126.

Setting Up Naming Parameters

When you configure your hardware iSCSI initiators, make sure their names and IP addresses are formatted properly.

See "Naming Requirements" on page 111.

To set up the iSCSI name, alias, and IP address for the hardware initiator

1 In the iSCSI Initiator Properties dialog box, click **Configure**.

2 To change the default iSCSI name for your initiator, enter the new iSCSI name.

You can use the default name that the vendor supplied. If you change the default name, make sure the new name you enter is properly formatted. Otherwise, some storage devices may not recognize the hardware iSCSI initiator.

3 Enter the iSCSI alias.

The alias is a name that you use to identify the hardware iSCSI initiator.

4 Enter all of the required values in the **Hardware Initiator Properties** group.

5 Click **OK** to save your changes.

6 Reboot the server for the changes to take effect.

Setting up Discovery Addresses for Hardware Initiators

Set up target discovery addresses so that the hardware initiator can determine which storage resource on the network is available for access. You can do this with either dynamic discovery or static discovery.

See "Discovery Methods" on page 112.

With dynamic discovery, a particular iSCSI server supplies a list of targets to your ESX Server 3 host.

To set up target discovery addresses using dynamic discovery

1 In the iSCSI Initiator Properties dialog box, click the **Dynamic Discovery** tab.

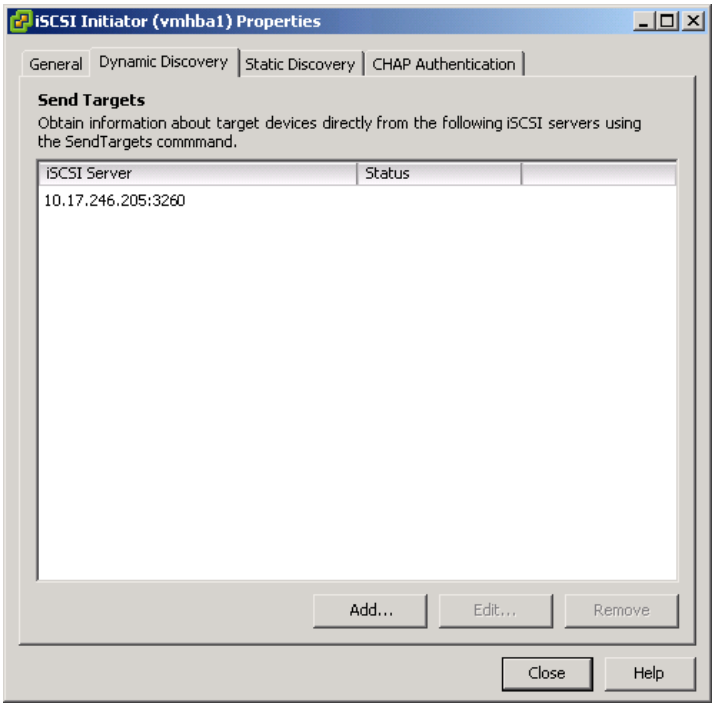

2 To add a new iSCSI server that your ESX Server 3 host can use for a dynamic discovery session, click **Add**.

3 In the Add Send Targets Server dialog box, enter the IP address of the iSCSI server and click **OK**.

After your ESX Server 3 host establishes the dynamic discovery session with this server, the server responds by providing a list of targets available to your ESX Server 3 host. The names and IP addresses of these targets appear on the **Static Discovery** tab.

4 To change the IP address of the iSCSI server or remove the server, select the IP address and click **Edit** or **Remove**.

With hardware initiators, in addition to the dynamic discovery method, you can also use static discovery, where you manually enter the IP addresses and the iSCSI names of the targets to be contacted.

To set up target discovery address by using static discovery

1 In the iSCSI Initiator Properties dialog box, click the **Static Discovery** tab.

If you previously used the dynamic discovery method, the tab displays any targets that the iSCSI server supplies to your ESX Server 3 host.

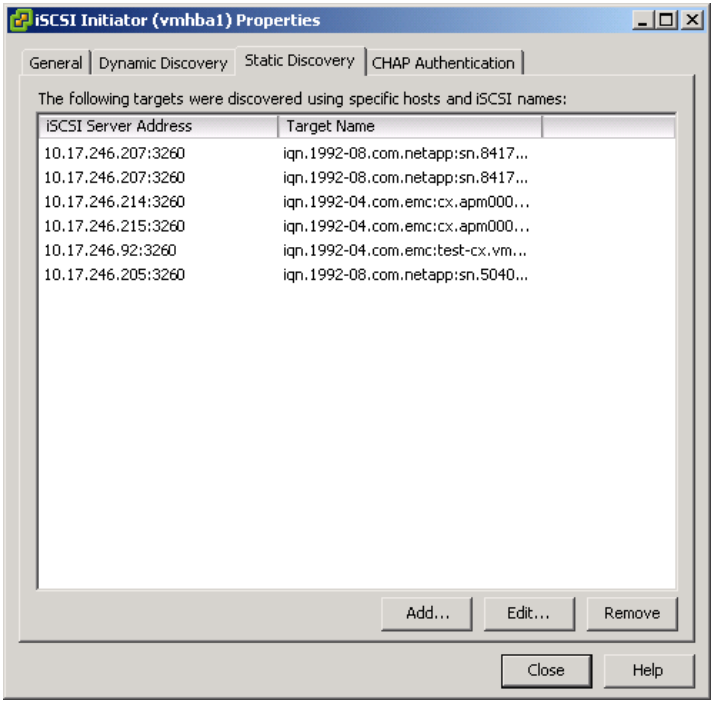

2 To add a target, click **Add** and enter the target's IP address and fully qualified name.

3 To change or delete a specific target, select the target and click **Edit** or **Remove**.

NOTE If you remove a target added by dynamic discovery, the target might be returned to the list the next time a rescan happens, the HBA is reset, or the system is rebooted.

Setting up CHAP Parameters for Hardware Initiators

When you configure your hardware iSCSI initiator, verify whether CHAP is enabled on the iSCSI storage. If it is enabled, enable it for your initiator, making sure that the CHAP authentication credentials match your iSCSI storage.

See "iSCSI Security" on page 112.

To set up CHAP parameters for the hardware initiator

1 In the iSCSI Initiator Properties dialog box, click the **CHAP Authentication** tab.

The tab displays the default CHAP parameters, if any.

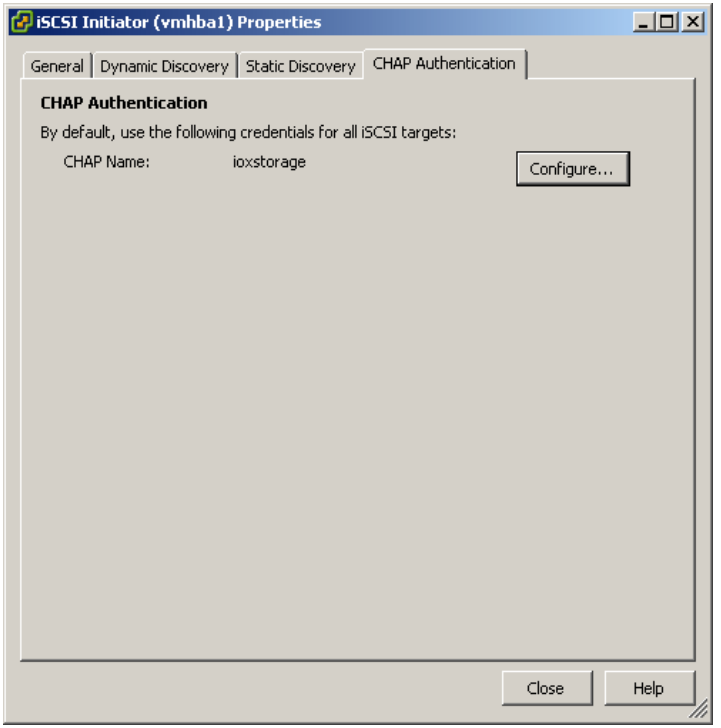

2 To make any changes to the existing CHAP parameters, click **Configure**.

3 To keep CHAP enabled, select **Use the following CHAP credentials**.

4 Either enter a new CHAP name, or select **Use initiator name**.

5 If needed, specify the CHAP secret.

All new targets will use the CHAP secret to authenticate the initiator.

6 Click **OK** to save changes.

NOTE If you disable CHAP, existing sessions remain until you reboot your ESX Server 3 host or the storage system forces a logout. After this, you can no longer connect to targets that require CHAP.

Adding iSCSI Storage Accessible Through Hardware Initiators

When you create a datastore on an iSCSI storage device accessible through a hardware initiator, the Add Storage wizard guides you through the configuration.

To create a datastore on a hardware iSCSI device

1 Log in to the VI Client and select a server from the inventory panel.

2 Click the **Configuration** tab, and click **Storage** in the **Hardware** panel.

3 Click **Add Storage**.

4 Select the **Disk/LUN** storage type and click **Next**.

5 Select the iSCSI device you want to use for your datastore, and click **Next**.

The Current Disk Layout page opens. If the disk you are formatting is blank, the Current Disk Layout page automatically presents the entire disk space to you for storage configuration.

6 If the disk is not blank, review the current disk layout in the top panel of the Current Disk Layout page and select a configuration option from the bottom panel:

- **Use the entire device** — Select this option to dedicate the entire disk or LUN to a single VMFS datastore. VMware recommends that you select this option.

WARNING If you select this option, any file systems or data previously stored on this device will be destroyed.

- **Use free space** — Select this option to deploy a VMFS datastore in the remaining free space of the disk.

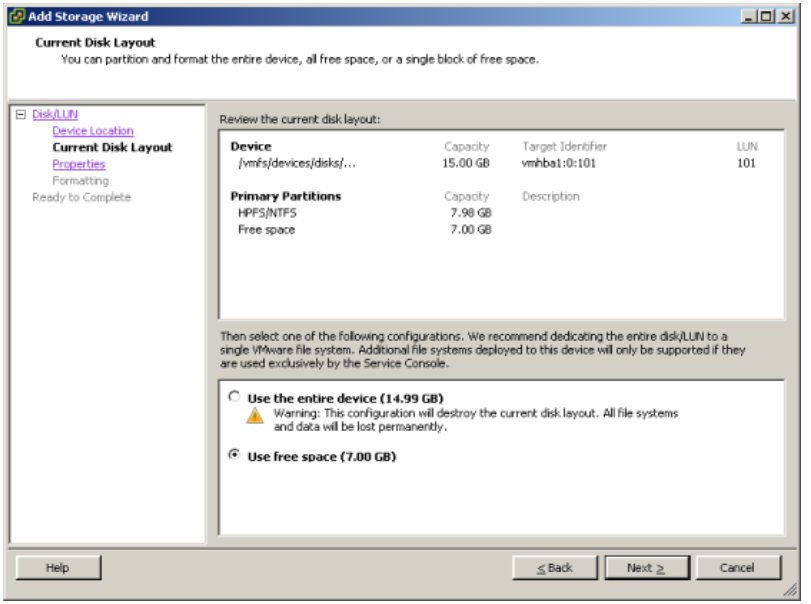

7 Click **Next**.

8 In the Disk/LUN–Properties page, enter a datastore name and click Next .

9 If needed, adjust the file system values and capacity you use for the datastore. By default, the entire free space available on the storage device is offered to you.

10 Click **Next**.

11 Review the datastore information, and click **Finish**.

This creates the datastore on the hardware-initiated iSCSI device.

12 Click **Refresh**.

Configuring Software iSCSI Initiators and Storage

With the software-based iSCSI implementation, you can use a standard network adapter to connect your ESX Server 3 system to a remote iSCSI target on the IP network. The ESX Server 3 software iSCSI initiator built into VMkernel facilitates this connection communicating with the network adapter through the network stack.

Before you configure datastores that use software initiators to access the iSCSI storage, enable network connectivity, then and configure the software iSCSI initiator.

Setting Up the iSCSI Storage Accessible Through Software Initiators

Perform the following tasks when you prepare and set up datastores that use software initiators to access the iSCSI storage device.

1 Create a VMkernel port to handle iSCSI networking.

  See "VMkernel Networking Configuration" on page 30 and "Networking Configuration for Software iSCSI Storage" on page 73.

2 Configure a service console connection for software iSCSI.

  See "Opening Firewall Ports for Supported Services and Management Agents" on page 188.

3 Configure the software iSCSI initiator.

  See "Configuring Software iSCSI Initiators" on page 123.

4 Rescan for new iSCSI LUNs.

  See "Performing a Rescan" on page 126.

5 Set up the datastore.

  See "Adding iSCSI Storage Accessible Through Software Initiators" on page 125.

Viewing Software iSCSI Initiators

The software iSCSI adapter that your ESX Server 3 system uses to access an iSCSI storage device appears on the list of available adapters. You can use the VI Client to review its properties.

To view the software iSCSI initiator properties

1 Log in to the VI Client, and select a server from the inventory panel.

2 Click the **Configuration** tab, and click **Storage Adapters** in the **Hardware** panel.

  The list of available storage adapters appears.

3 Under iSCSI Software Adapter, choose the available software initiator.

If the initiator is enabled, the Details panel displays the initiator's model, IP address, iSCSI name, discovery methods, iSCSI alias, and any discovered targets.

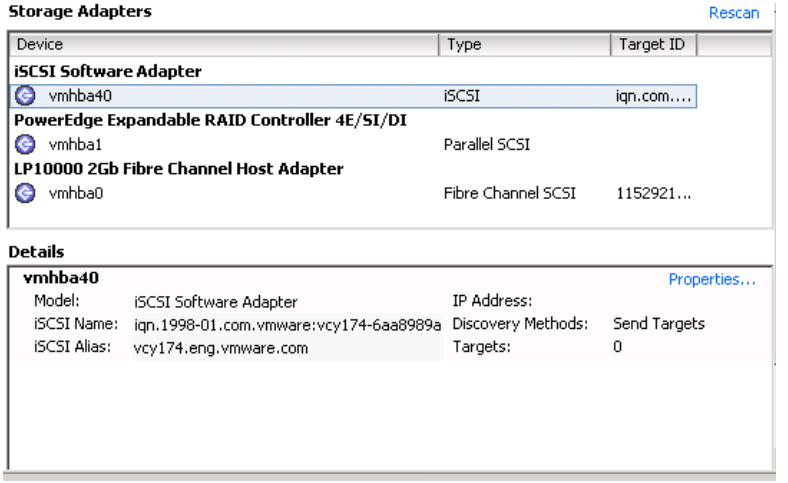

4 Click **Properties**.

The iSCSI Initiator Properties dialog box opens. The **General** tab displays additional characteristics of the software initiator.

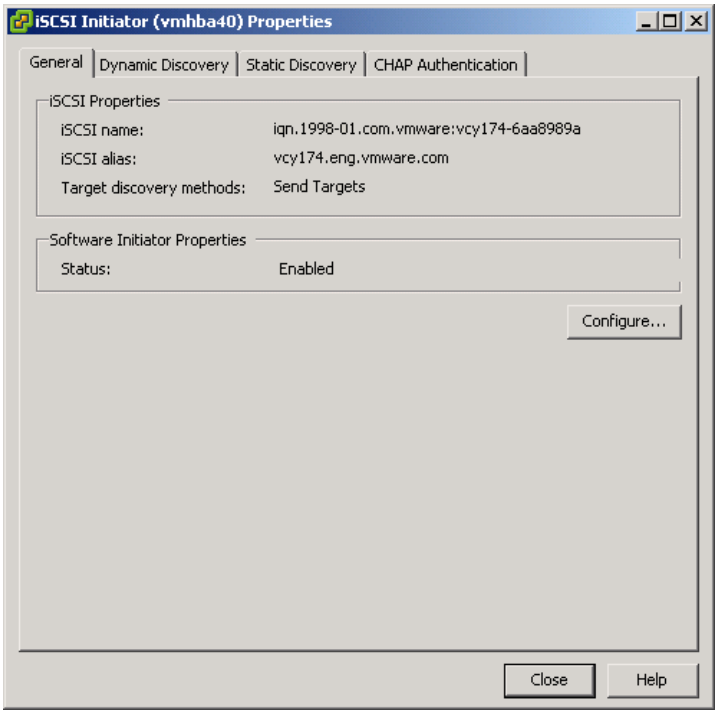

You can now configure your software initiator or change its default characteristics.

Configuring Software iSCSI Initiators

When you configure the software iSCSI initiator, you enable your initiator and set up its target addresses. VMware also recommends that you set up the CHAP parameters. After you configure your software iSCSI initiator, perform a rescan, so that all LUNs that the initiator has access to appear on the list of storage devices available to your ESX Server 3 system. See "Performing a Rescan" on page 126.

Enabling Software iSCSI Initiators

Enable your software iSCSI initiator so that ESX Server 3 can use it.

To enable the software iSCSI initiator

1 In the iSCSI Initiator Properties dialog box, click **Configure**.

2 To enable the initiator, select **Enabled**.

3 To change the default iSCSI name for your initiator, enter the new name.

Make sure the name you enter is properly formatted. Otherwise, some storage devices might not recognize the software iSCSI initiator. See "Naming Requirements" on page 111.

4 Click **OK** to save your changes.

Setting Up Discovery Addresses for Software Initiators

Set up target discovery addresses so that the software initiator can determine which storage resource on the network is available for access.

NOTE With software initiators, only the dynamic discovery method is available.

See "Discovery Methods" on page 112.

To set up target discovery addresses for the software initiator

1 In the iSCSI Initiator Properties dialog box, click the **Dynamic Discovery** tab.

2 To add a new iSCSI target that your ESX Server 3 host can use for a dynamic discovery session, click **Add**.

3 Enter the Send Targets server IP address and click **OK**.

4 To change or delete the Send Targets server, select the server and click **Edit** or **Remove**.

Setting Up CHAP Parameters for Software Initiators

When you configure your software iSCSI initiator, verify whether CHAP is enabled on the iSCSI storage. If it is enabled, enable it for the initiator, making sure that the CHAP authentication credentials match your iSCSI storage.

See "iSCSI Security" on page 112.

To set up CHAP parameters for the software initiator

1 In the iSCSI Initiator Properties dialog box, click the **CHAP Authentication** tab.

2 To specify CHAP parameters, click **Configure**.

3 To keep CHAP enabled, select **Use the following CHAP credentials**.

4 Enter a CHAP name, or select **Use initiator name**.

5 If needed, specify the CHAP secret.

All new targets will use the CHAP secret to authenticate the initiator.

6 Click **OK** to save changes.

NOTE If you disable CHAP, existing sessions remain until you reboot your ESX Server 3 host or the storage system forces a logout. After this, you can no longer connect to targets that require CHAP.

Adding iSCSI Storage Accessible Through Software Initiators

When you create a datastore on an iSCSI storage device accessible through software initiators, the Add Storage wizard guides you through the configuration.

To create a datastore on an iSCSI device accessed through software initiators

1 Log in to the VI Client, and select a server from the inventory panel.

2 Click the **Configuration** tab and click **Storage** in the **Hardware** panel.

3 Click **Add Storage**.

4 Select the **Disk/LUN** storage type and click **Next**.

5 Select the iSCSI device to use for your datastore, and click **Next**.

The Current Disk Layout page opens. If the disk you are formatting is blank, the Current Disk Layout page automatically presents the entire disk space to you for storage configuration.

6 If the disk is not blank, review the current disk layout in the top panel of the Current Disk Layout page and select a configuration option from the bottom panel:

- **Use the entire device** — Select this option to dedicate the entire disk or LUN to a single VMFS datastore. VMware recommends that you select this option.

WARNING If you select this option, any file systems or data previously stored on this device will be destroyed.

- **Use free space** — Select this option to deploy a VMFS datastore in the remaining free space of the disk.

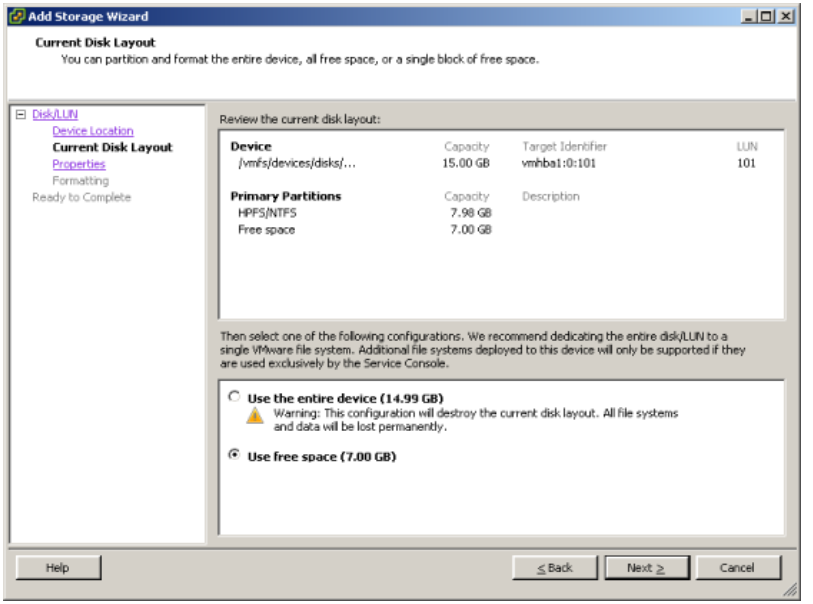

7 Click **Next**.

8 In the Disk/LUN–Properties page, enter a datastore name and click Next.

9 If needed, adjust the file system values and capacity you use for the datastore. By default, the entire free space available on the storage device is offered to you.

10 Click **Next**.

11 Review the datastore information, and click **Finish**.

This process creates the datastore on the software iSCSI storage device.

12 Click **Refresh**.

Performing a Rescan

Perform a rescan if any of the following events occur:

- When you make changes to storage disks or LUNs available to your ESX Server 3 system.
- When you make changes to storage adapters.

- When you create a new datastore or remove an existing one.
- When you reconfigure an existing datastore, for example, when you add a new extent.

NOTE After you mask all paths to a LUN, rescan all adapters with paths to the LUN in order to update the configuration.

To perform a rescan

1 In the VI Client, select a host, and click the **Configuration** tab.

2 Choose **Storage Adapters** in the Hardware panel and click **Rescan** above the Storage Adapters panel.

NOTE You can also right-click an individual adapter and click **Rescan** to rescan just that adapter.

3 To discover new disks or LUNs, select **Scan for New Storage Devices**.

If new LUNs are discovered, they appear in the disk/LUN list.

4 To discover new datastores or update a datastore after its configuration has been changed, select **Scan for New VMFS Volumes**.

If new datastores or VMFS volumes are discovered, they appear in the datastore list.

Network Attached Storage

This section contains information about network attached storage (NAS).

ESX Server 3 supports using NAS through the NFS protocol.

How Virtual Machines Use NFS

The NFS protocol that ESX Server 3 supports enables communication between an NFS client and NFS server. The client issues requests for information from the server, which replies with the result.

The NFS client built into ESX Server 3 lets you access the NFS server and use NFS volumes for storage. ESX Server 3 supports NFS Version 3 over TCP only.

You use the VI Client to configure NFS volumes as datastores. Configured NFS datastores appear in the VI Client and you can use them to store virtual disk files in the same way you use VMFS-based datastores.

The virtual disks that you create on NFS-based datastores use a disk format dictated by the NFS server, typically a thin-disk format that requires on-demand space allocation. If the virtual machine runs out of space while writing to this disk, the VI Client notifies you that more space is needed. You have the following options:

- Free up additional space on the volume, so that the virtual machine continues writing to the disk.
- Terminate the virtual machine session. Terminating the session shuts down the virtual machine.

Figure 6-4 depicts a virtual machine using the NFS volume to store its files.

Figure 6-4. NFS Storage

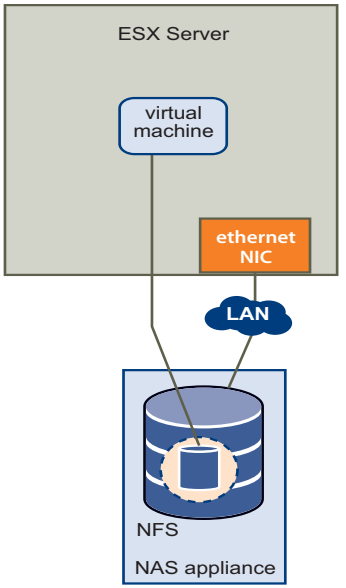

In this configuration, ESX Server 3 connects to the NFS server, which stores the virtual disk files.

WARNING When ESX Server 3 accesses a virtual machine disk file on an NFS-based datastore, a special `.lck–XXX` lock file is generated in the same directory where the disk file resides to prevent other ESX Server 3 hosts from accessing this virtual disk file. Do not remove the `.lck–XXX` lock file, because without it, the running virtual machine cannot access its virtual disk file.

NFS Volumes and Virtual Machine Delegate Users

If you are planning to create, configure, or administer virtual machines on an NFS-based datastore, assign NFS access privileges to a special user, known as the delegate user.

By default, the delegate user for the ESX Server 3 host is `root`. However, having `root` as the delegate user might not work for all NFS volumes. In some cases, to protect NFS volumes from unauthorized access, NFS administrators export the volumes with the `root squash` option turned on. When `root squash` is on, the NFS server treats access by the root user as access by any unprivileged user and might refuse the ESX Server 3 host access to virtual machine files stored on the NFS volume.

You can change the delegate user to a different identity through experimental ESX Server 3 functionality. This identity must match the owner of the directory on the NFS server, or the ESX Server 3 host cannot perform file level operations.

See "Virtual Machine Delegates for NFS Storage" on page 232.

CAUTION Changing the delegate user for an ESX Server 3 host is experimental and, currently, VMware provides limited support for this feature.

Configuring ESX Server 3 to Access NFS Volumes

NFS requires network connectivity to access data stored on remote servers. Before configuring NFS, you must first configure networking for VMotion and IP storage.

For information on configuring a network, see "VMkernel Networking Configuration" on page 30.

Creating an NFS-Based Datastore

When you create a datastore on an NFS volume, the Add Storage wizard guides you through the configuration steps.

To mount an NFS volume

1 Log in to the VI Client, and select a server from the inventory panel.

2 Click the **Configuration** tab, and click **Storage** in the **Hardware** panel.

3 Click **Add Storage**.

4 Select **Network File System** as the storage type and click **Next**.

5 Enter the server name, the mount point folder name, and the datastore name.

6 Click **Next**.

7 In the Network File System Summary page, review the configuration options and click **Finish**.

Creating a Diagnostic Partition

To run successfully, your ESX Server 3 needs to have a diagnostic partition, or a dump partition, to store core dumps for debugging and technical support. You can create the diagnostic partition on a local disk, or on a private or shared SAN LUN.

A diagnostic partition cannot be located on an iSCSI LUN accessed through a software initiator.

Each ESX Server 3 host must have a diagnostic partition of 100MB. If multiple ESX Server 3 hosts share a SAN, configure a diagnostic partition with 100MB for each host.

NOTE If you selected **Recommended Partitioning** when you installed ESX Server 3, the installer automatically created a diagnostic partition for your host. The **Diagnostic** option doesn't appear on the Select Storage Type page. If you selected **Advanced Partitioning** and chose not to specify the diagnostic partition during installation, configure it now. For information on ESX Server 3 partitioning, see the *Installation Guide*.

To create a diagnostic partition

1 Log in to the VI Client, and select a server from the inventory panel.

2 Click the **Configuration** tab, and click **Storage** in the **Hardware** panel.

3 Click **Add Storage**.

The Select Storage Type page appears.

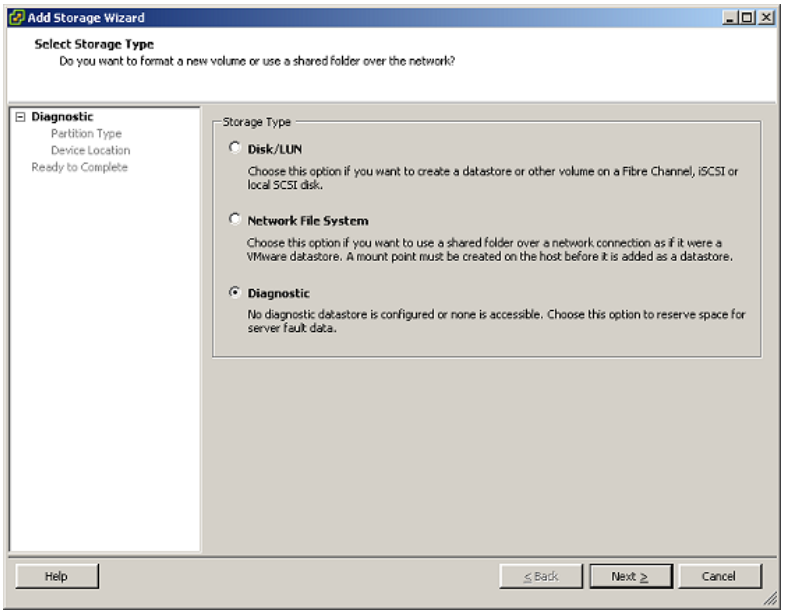

4 Select **Diagnostic** and click **Next**.

If you don't see **Diagnostic** as an option, ESX Server 3 host already has a diagnostic partition. You can query and scan your host's diagnostic partition by using the `esxcfg-dumppart` service console command. See "ESX Server 3 Technical Support Commands" on page 277.

5 Specify the type of the diagnostic partition:

- **Private Local** – Creates the diagnostic partition on a local disk. This partition stores fault information only for your ESX Server 3 host.
- **Private SAN Storage** – Creates the diagnostic partition on a non-shared SAN LUN. This partition stores fault information only for your ESX Server 3 host.
- **Shared SAN Storage** – Creates the diagnostic partition on a shared SAN LUN. This partition is accessed by multiple hosts and may store fault information for more than one host.

Click **Next**.

6 Select the device to use for your diagnostic partition and click **Next**.

7 Review the partition configuration information and click **Finish**.

7 Managing Storage

This chapter contains information about managing existing datastores and file systems that comprise datastores. The chapter discusses the following topics:

- "Managing Datastores" on page 133
- "Editing VMFS Datastores" on page 134
- "Managing Multiple Paths" on page 137
- "The vmkfstools Commands" on page 144

Managing Datastores

An ESX Server 3 system uses datastores to store all files associated with its virtual machines. The datastore is a logical storage unit, which can use disk space on one physical device, one disk partition, or span several physical devices. The datastore can exist on different types of physical devices, including SCSI, iSCSI, Fibre Channel SAN, or NFS.

NOTE As an alternative to using the datastore, your virtual machine can directly access raw devices by using a mapping file (RDM) as a proxy. See "Raw Device Mapping Characteristics" on page 150.

For more information on datastores, see "Datastores" on page 91.

Datastores are added to the VI Client in one of the following ways:

- Discovered when a host is added to the inventory. When you add a host to the inventory, the VI Client displays any datastores that the host can recognize.
- Created on an available storage device. You can use the **Add Storage** command to create and configure a new datastore.

After you create the datastores, you can use them to store virtual machine files. When needed, you can modify the datastores. For example, you can add extents to your datastore, rename, or remove it.

You can remove a datastore that you don't use.

CAUTION Removing a datastore from the ESX Server 3 system breaks the connection between the system and the storage device that holds the datastore and stops all functions of that storage device.

You cannot remove a datastore if it holds virtual disks of a currently running virtual machine.

To remove a datastore

1 Log in to the VI Client and select a server from the inventory panel.

2 Click the **Configuration** tab and click **Storage**.

3 Select the datastore to remove and click **Remove**.

4 Confirm that you want to remove the datastore.

5 Perform a rescan on all servers that see the datastore.

Editing VMFS Datastores

Datastores that use the VMFS format are deployed on SCSI-based storage devices.

After you create a VMFS-based datastore, you can modify it by renaming or expanding it. If you have any VMFS-2 datastores, you can upgrade them to VMFS-3 format.

Upgrading Datastores

ESX Server 3 includes VMFS version 3 (VMFS-3). If your datastore was formatted with VMFS-2, you can read files stored on VMFS-2, but you cannot use them. To use the files, upgrade VMFS-2 to VMFS-3.

When you upgrade VMFS-2 to VMFS-3, the ESX Server 3 file-locking mechanism ensures that no remote ESX Server 3 or local process is accessing the VMFS volume being converted. ESX Server 3 preserves all files on the datastore.

As a precaution, before you use the upgrade option, consider the following:

- Commit or discard any changes to virtual disks in the VMFS-2 volume you plan to upgrade.
- Back up the VMFS-2 volume you want to upgrade.
- Be sure that no powered-on virtual machines are using this VMFS-2 volume.
- Be sure that no other ESX Server host is accessing this VMFS-2 volume.

CAUTION The VMFS-2 to VMFS-3 conversion is a one-way process. After you convert the VMFS-based datastore to VMFS-3, you cannot revert it back to VMFS-2.

To be able to upgrade the VMFS-2 file system, its file block size should not exceed 8MB.

To upgrade the VMFS-2 to VMFS-3

1 Log in to the VI Client and select a server from the inventory panel.

2 Click the **Configuration** tab and click **Storage**.

3 Select the datastore that uses the VMFS-2 format.

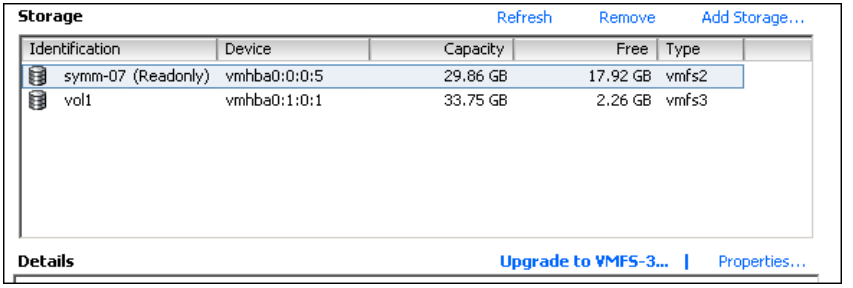

4 Click **Upgrade to VMFS-3**.

5 Perform a rescan on all hosts that see the datastore.

Changing the Names of Datastores

You can change the name of an existing VMFS-based datastore.

To edit the name of the datastore

1 Log in to the VI Client and select a server from the inventory panel.

2 Click the **Configuration** tab and click **Storage**.

3 Select the datastore whose name you want to edit and click **Properties**.

4 In the General panel, click **Change**.

5 Enter the new datastore name and click **OK**.

Adding Extents to Datastores

You can expand a datastore that uses the VMFS format by attaching a hard-disk partition as an extent. The datastore can span 32 physical storage extents.

You can dynamically add the new extents to the datastore when you need to create new virtual machines on this datastore, or when the virtual machines running on this datastore require more space.

To add one or more extents to the datastore

1 Log in to the VI Client and select a server from the inventory panel.

2 Click the **Configuration** tab and click **Storage**.

3 Select the datastore to expand and click **Properties**.

4 In the Extents panel, click **Add Extent**.

5 Select the disk to add as the new extent and click **Next**.

6 Review the current layout of the disk you are using for the extent to make sure the disk doesn't contain any important information.

CAUTION If a disk or partition you add was formatted previously, it will be reformatted and lose the file systems and any data it contains.

7 Set the capacity for the extent.

By default, the entire free space available on the storage device is offered to you.

8 Click **Next**.

9 Review the proposed extent layout and the new configuration of your datastore, and click **Finish**.

10 Perform a rescan on all servers that see the datastore.

Managing Multiple Paths

To maintain a constant connection between the ESX Server 3 host and its direct attached or networked storage, ESX Server 3 supports multipathing. *Multipathing* is a technique that lets you use more than one physical element on a path responsible for transferring data between the ESX Server 3 host and the external storage device. In case of a failure of any element on the path, an HBA, switch, storage processor (SP), or cable, ESX Server 3 can use a redundant path. The process of detecting a failed path and switching to another path is called *path failover*. This use of failover paths helps ensure uninterrupted traffic between the ESX Server 3 system and storage devices. To support multipathing, ESX Server 3 does not require specific failover drivers.

NOTE A virtual machine fails in an unpredictable way if all paths to the storage device where you stored your virtual machine disks become unavailable.

By default, the ESX Server 3 host uses only one path, called the *active path*, to communicate with a particular storage device at any given time.

When selecting the active path, ESX Server 3 follows these multipathing policies:

- **Most Recently Used** — As its active path, the ESX Server 3 host selects the path it used most recently. If this path becomes unavailable, the host switches to an alternative path and continues to use the new path as the active path.

  The Most Recently Used policy is required for *active/passive* storage arrays, in which one storage processor remains passive waiting for the other to fail.

- **Fixed** — The ESX Server 3 host always uses the designated preferred path to the storage device as the active path. If the ESX Server 3 host cannot access the storage through the preferred path, it tries the alternative path, which then becomes the active path. The host automatically reverts back to the preferred path as soon as it is available.

  VMware recommends the Fixed policy for *active/active* storage arrays, in which all storage processors can pass the storage traffic and all paths can be active at all times, unless a path fails. Most iSCSI storage systems are active/active.

  NOTE VMware recommends that you do not manually change **Most Recently Used** to **Fixed**. The system automatically sets this policy for those arrays that require it.

- **Round Robin** – The ESX Server 3 host uses an automatic path selection rotating through all available paths. In addition to path failover, round robin supports load balancing across the paths.

  In this release, round robin load balancing is experimental and not supported for production use. See the *Round-Robin Load Balancing* white paper.

Multipathing with Local Storage and Fibre Channel SANs

In a simplest multipathing local storage topology, you can use one ESX Server 3 host, which has two HBAs. The ESX Server 3 host connects to a dual-port local storage system through two cables. Using this configuration, you can ensure fault tolerance if one of the connection elements between the ESX Server 3 host and local storage system fails.

To support path switching with FC SAN, the ESX Server 3 host typically has two or more HBAs available, from which the storage array can be reached using one or more switches. Alternatively, the setup could include one HBA and two storage processors so that the HBA can use a different path to reach the disk array.

In Figure 7-1, multiple paths connect each server with the storage device. For example, if HBA1 or the link between HBA1 and the switch fails, HBA2 takes over and provides the connection between the server and the switch. The process of one HBA taking over for another is called HBA failover.

Figure 7-1. Fibre Channel Multipathing

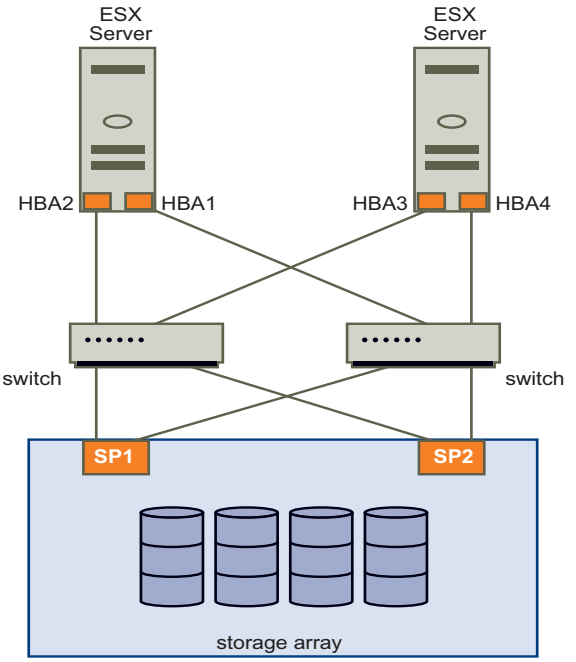

Similarly, if SP1 or the link between SP1 and the switch breaks, SP2 takes over and provides the connection between the switch and the storage device. This process is called SP failover. VMware ESX Server 3 supports HBA and SP failover with its multipathing capability.

See the *Fibre Channel SAN Configuration Guide*.

Multipathing with iSCSI SAN

With iSCSI storage, ESX Server 3 takes advantage of the multipathing support built into the IP network, which allows the network to perform routing, as Figure 7-2 shows. Through Dynamic Discovery, iSCSI initiators obtain a list of target addresses that the initiators can use as multiple paths to iSCSI LUNs for failover purposes.

Figure 7-2. iSCSI Multipathing

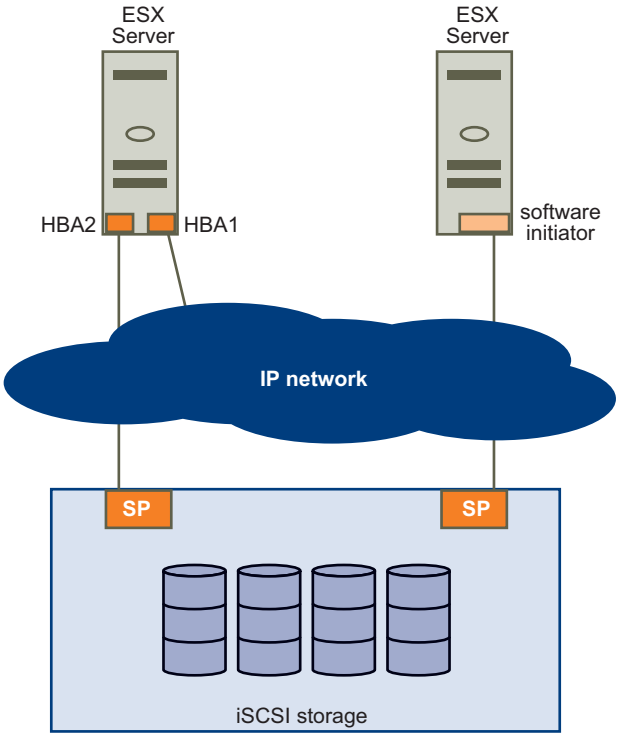

In addition, with the software-initiated iSCSI, you can use NIC teaming, so that multipathing is performed through the networking layer in the VMkernel. See "Networking" on page 17.

See the *iSCSI SAN Configuration Guide*.

Viewing the Current Multipathing Status

Use the VI Client to view the current multipathing state.

To view the current multipathing state

1 Log in to the VI Client and select a server from the inventory panel.

2 Click the **Configuration** tab and click **Storage** in the Hardware panel.

3 From the list of configured datastores, select the datastore whose paths you want to view or configure.

The Details panel shows the total number of paths being used to access the datastore and whether any of them are broken or disabled.

4 Click **Properties**.

The Volume Properties dialog box for the selected datastore opens.

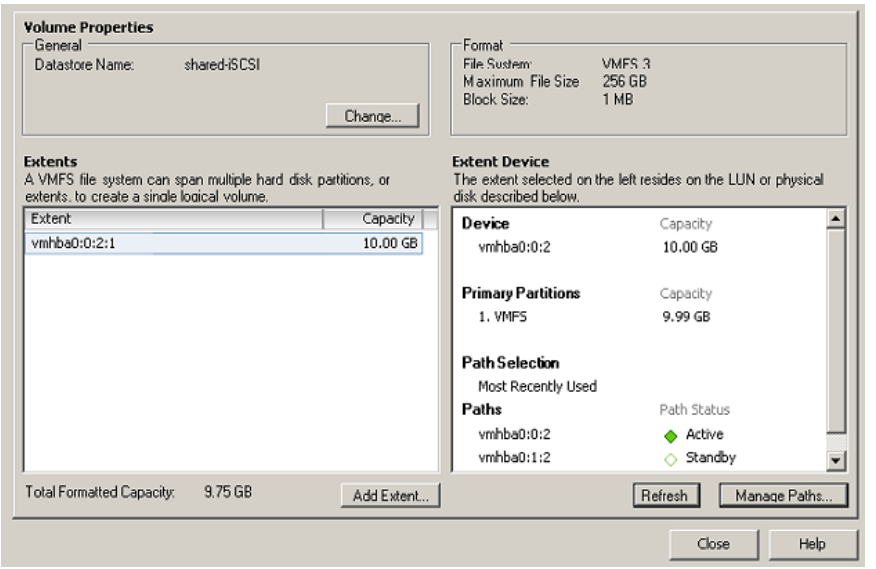

The Extent Device panel includes information on the multipathing policy that the ESX Server 3 host uses to access the datastore and on the status of each path. The following path information can appear:

- **Active** – The path is working and is the current path being used for transferring data.
- **Disabled** – The path is disabled and no data can be transferred.
- **Standby** – The path is working but is not currently being used for transferring data.
- **Broken** – The software cannot connect to the disk through this path.

5 Click **Manage Paths** to open the Manage Paths dialog box.

If you are using the **Fixed** path policy, you can see which path is the preferred path. The preferred path is marked with an asterisk (*) in the fourth column.

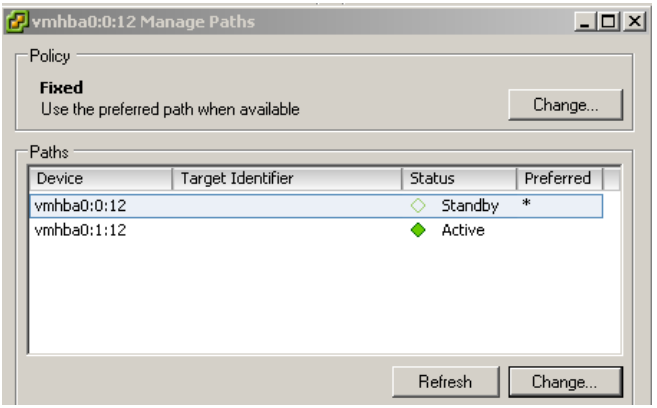

You can use the Manage Paths dialog box to enable or disable your paths, set multipathing policy, and specify the preferred path.

Setting Multipathing Policies for LUNs

Use the Manage Paths dialog box to set multipathing policy and specify the preferred path for the Fixed policy. If you are managing paths for RDMs, see "To manage paths" on page 158.

The Manage Paths dialog box shows the list of different paths to the disk, with the multipathing policy for the disk and the connection status for each path. It also shows the preferred path to the disk.

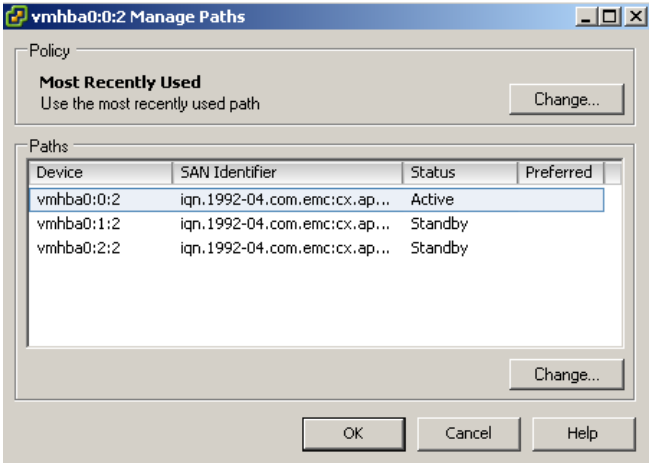

To set the multipathing policy

1 In the Policy panel, click **Change**.

2 Select one of the following options:

- **Fixed**
- **Most Recently Used**
- Round Robin (Experimental)

3 Click **OK** and click **Close** to save your settings and return to the Configuration page.

NOTE VMware recommends **Most Recently Used** for active/passive storage devices.

If you set the path policy to **Fixed**, specify the preferred path that the host should use when it is available.

To set the preferred path

1 In the Paths panel, select the path to make the preferred path, and click **Change**.

2 In the Preference pane, click **Preferred**.

If **Preferred** does not appear as an option, make sure that the **Path Policy** is **Fixed**.

3 Click **OK** twice to save your settings and exit the dialog boxes.

Disabling Paths

To temporarily disable paths for maintenance or any other reasons, use the VI Client.

To disable a path

1 In the Paths panel, select the path to disable and click **Change**.

2 Select **Disabled** to disable the path.

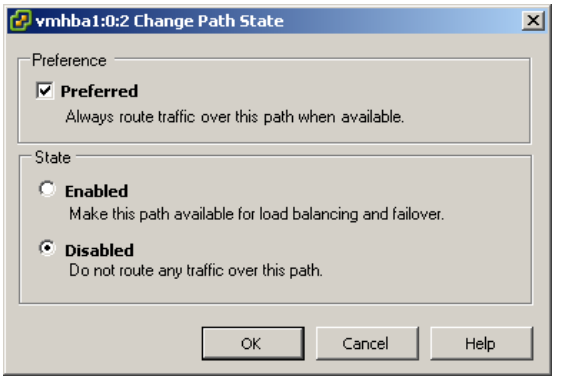

3 Click **OK** twice to save your settings and exit the dialog boxes.

The vmkfstools Commands

In addition to using VI Client, you can use the `vmkfstools` program to manage physical storage devices and to create and manipulate VMFS datastores and volumes on your ESX Server 3 host.

For a list of supported `vmkfstools` commands, see "Using vmkfstools" on page 283.

8 Raw Device Mapping

Raw device mapping (RDM) provides a mechanism for a virtual machine to have direct access to a LUN on the physical storage subsystem (Fibre Channel or iSCSI only). This chapter contains information about RDMs.

This chapter discusses the following topics:

- "About Raw Device Mapping" on page 145
- "Raw Device Mapping Characteristics" on page 150
- "Managing Mapped LUNs" on page 155

About Raw Device Mapping

RDM is a mapping file in a separate VMFS volume that acts as a proxy for a raw physical device, a SCSI device used directly by a virtual machine. The RDM contains metadata used to manage and redirect disk accesses to the physical device. The file gives you some of the advantages of direct access to a physical device while keeping some advantages of a virtual disk in the VMFS file system. As a result, it merges VMFS manageability with raw device access.

RDMs can be described in terms such as "Mapping a raw device into a datastore," "mapping a system LUN", or "mapping a disk file to a physical disk volume." All these terms refer to RDMs.

Figure 8-1. Raw Device Mapping

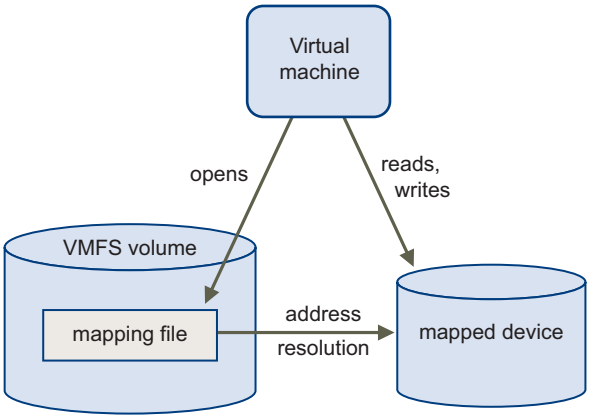

Although VMware recommends that you use VMFS datastores for most virtual disk storage, on certain occasions, you might need to use raw LUNs, or logical disks located in a SAN.

For example, you need to use raw LUNs with RDMs in the following situations:

- When SAN snapshot or other layered applications are run in the virtual machine. The RDM better enables scalable backup offloading systems by using features inherent to the SAN.
- In any MSCS clustering scenario that spans physical hosts — virtual-to-virtual clusters as well as physical-to-virtual clusters. In this case, cluster data and quorum disks should be configured as RDMs rather than as files on a shared VMFS.

Think of an RDM as a symbolic link from a VMFS volume to a raw LUN (see Figure 8-1). The mapping makes LUNs appear as files in a VMFS volume. The RDM, not the raw LUN, is referenced in the virtual machine configuration. The RDM contains a reference to the raw LUN.

Using RDMs, you can:

- Use VMotion to migrate virtual machines using raw LUNs.
- Add raw LUNs to virtual machines using the VI Client.
- Use file system features such as distributed file locking, permissions, and naming.

Two compatibility modes are available for RDMs:

- Virtual compatibility mode allows an RDM to act exactly like a virtual disk file, including the use of snapshots.
- Physical compatibility mode allows direct access of the SCSI device, for those applications that need lower level control.

Benefits of Raw Device Mapping

An RDM provides a number of benefits, but it should not be used in every situation. In general, virtual disk files are preferable to RDMs for manageability. However, when you need raw devices, you must use the RDM. The following list highlights the benefits of the RDM:

- **User-Friendly Persistent Names** – Provides a user-friendly name for a mapped device. When you use an RDM, you do not need to refer to the device by its device name. You refer to it by the name of the mapping file, for example:

  `/vmfs/volumes/myVolume/myVMDirectory/myRawDisk.vmdk`

- **Dynamic Name Resolution** – Stores unique identification information for each mapped device. The VMFS file system associates each RDM with its current SCSI device, regardless of changes in the physical configuration of the server because of adapter hardware changes, path changes, device relocation, and so on.
- **Distributed File Locking** – Makes it possible to use VMFS distributed locking for raw SCSI devices. Distributed locking on an RDM makes it safe to use a shared raw LUN without losing data when two virtual machines on different servers try to access the same LUN.
- **File Permissions** – Makes file permissions possible. The permissions of the mapping file are enforced at file-open time to protect the mapped volume.
- **File System Operations** – Makes it possible to use file system utilities to work with a mapped volume, using the mapping file as a proxy. Most operations that are valid for an ordinary file can be applied to the mapping file and are redirected to operate on the mapped device.
- **Snapshots** – Makes it possible to use virtual machine snapshots on a mapped volume.

NOTE Snapshots are not available when the RDM is used in physical compatibility mode.

- **VMotion** – Lets you migrate a virtual machine with VMotion. The mapping file acts as a proxy to allow VirtualCenter to migrate the virtual machine by using the same mechanism that exists for migrating virtual disk files. See Figure 8-2.

Figure 8-2. VMotion of a Virtual Machine Using Raw Device Mapping

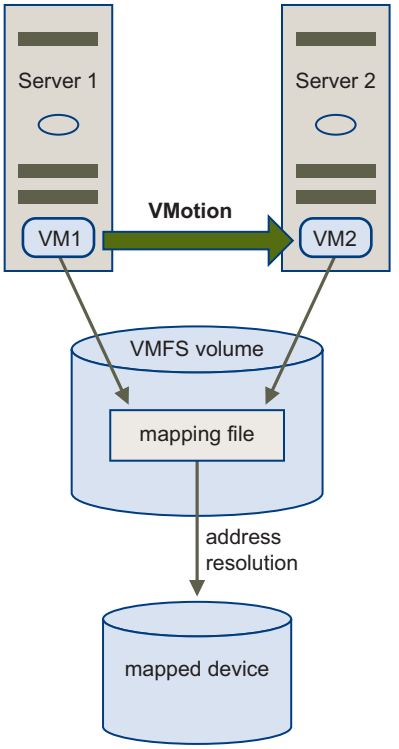

- **SAN Management Agents** – Makes it possible to run some SAN management agents inside a virtual machine. Similarly, any software that needs to access a device by using hardware-specific SCSI commands can be run in a virtual machine. This kind of software is called SCSI target-based software.

NOTE When you use SAN management agents, select a physical compatibility mode for the RDM.

- **N-Port ID Virtualization (NPIV)** – Makes it possible to use the NPIV technology that allows a single Fibre Channel HBA port to register with the Fibre Channel fabric using several worldwide port names (WWPNs). This ability makes the HBA port appear as multiple virtual ports, each having its own ID and virtual port name. Virtual machines can then claim each of these virtual ports and use them for all RDM traffic.

  NOTE NPIV can be used only for virtual machines with RDM disks.

  See the *Fibre Channel SAN Configuration Guide*.

VMware works with vendors of storage management software to ensure that their software functions correctly in environments that include ESX Server 3. Some applications of this kind are:

- SAN management software
- Storage resource management (SRM) software
- Snapshot software
- Replication software

Such software uses a physical compatibility mode for RDMs, so that the software can access SCSI devices directly.

Various management products are best run centrally (not on the ESX Server 3 machine), while others run well on the service console or on the virtual machines. VMware does not certify these applications or provide a compatibility matrix. To find out whether a SAN management application is supported in an ESX Server 3 environment, contact the SAN management software provider.

Limitations of Raw Device Mapping

When you plan to use an RDM, consider the following:

- **Not Available for Block Devices or Certain RAID Devices** – RDM (in the current implementation) uses a SCSI serial number to identify the mapped device. Because block devices and some direct-attach RAID devices do not export serial numbers, they cannot be used with RDMs.
- **Available with VMFS-2 and VMFS-3 Volumes Only** – RDM requires the VMFS-2 or VMFS-3 format. In ESX Server 3, the VMFS-2 file system is read-only. Upgrade it to VMFS-3 to use the files that VMFS-2 stores.

- **No Snapshots in Physical Compatibility Mode** – If you are using an RDM in physical compatibility mode, you cannot use a snapshot with the disk. Physical compatibility mode allows the virtual machine to manage its own snapshot or mirroring operations.

  Snapshots are available, however, in virtual mode. See "Virtual Compatibility Mode Compared to Physical Compatibility Mode" on page 151.

- **No Partition Mapping** – RDM requires the mapped device to be a whole LUN. Mapping to a partition is not supported.

Raw Device Mapping Characteristics

An RDM is a special mapping file in a VMFS volume that manages metadata for its mapped device. The mapping file is presented to the management software as an ordinary disk file, available for the usual file-system operations. To the virtual machine, the storage virtualization layer presents the mapped device as a virtual SCSI device.

Key contents of the metadata in the mapping file include the location of the mapped device (name resolution) and the locking state of the mapped device.

Figure 8-3. Mapping File Metadata

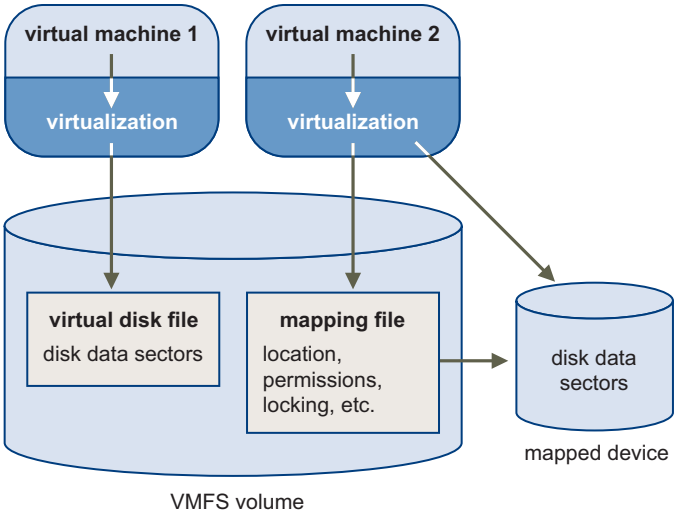

Virtual Compatibility Mode Compared to Physical Compatibility Mode

Virtual mode for an RDM specifies full virtualization of the mapped device. It appears to the guest operating system exactly the same as a virtual disk file in a VMFS volume. The real hardware characteristics are hidden. Virtual mode allows customers using raw disks to realize the benefits of VMFS such as advanced file locking for data protection and snapshots for streamlining development processes. Virtual mode is also more portable across storage hardware than physical mode, presenting the same behavior as a virtual disk file.

Physical mode for the RDM specifies minimal SCSI virtualization of the mapped device, allowing the greatest flexibility for SAN management software. In physical mode, the VMkernel passes all SCSI commands to the device, with one exception: the REPORT LUNs command is virtualized, so that the VMkernel can isolate the LUN for the owning virtual machine. Otherwise, all physical characteristics of the underlying hardware are exposed. Physical mode is useful to run SAN management agents or other SCSI target based software in the virtual machine. Physical mode also allows virtual-to-physical clustering for cost-effective high availability.

Figure 8-4. Virtual And Physical Compatibility Modes

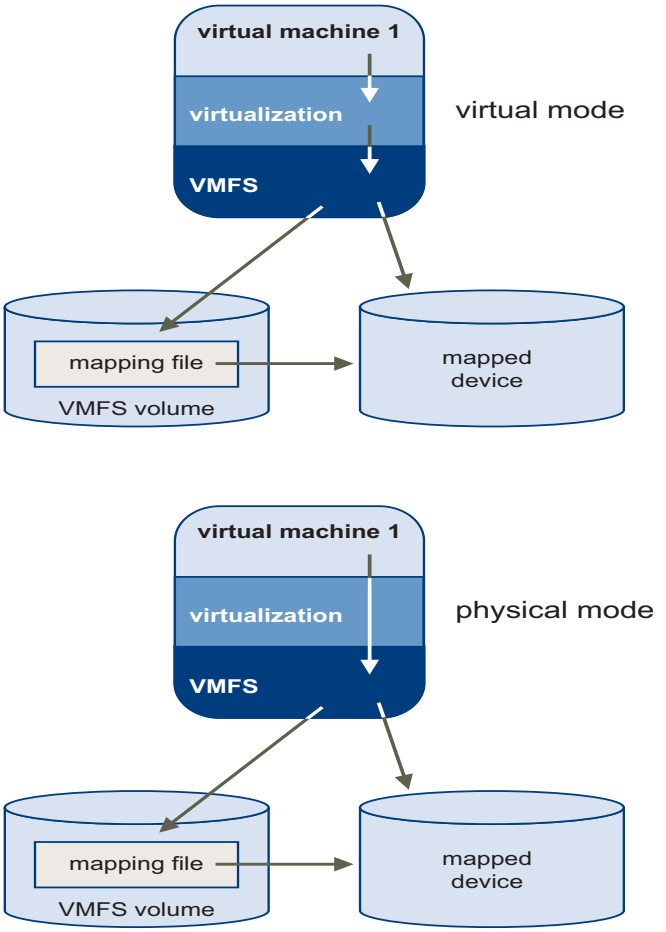

Dynamic Name Resolution

RDM lets you give a permanent name to a device by referring to the name of the mapping file in the `/vmfs` subtree.

The example in Figure 8-5 shows three LUNs. LUN 1 is accessed by its device name, which is relative to the first visible LUN. LUN 2 is a mapped device, managed by an RDM on LUN 3. The RDM is accessed by its path name in the `/vmfs` subtree, which is fixed.

Figure 8-5. Example of Name Resolution

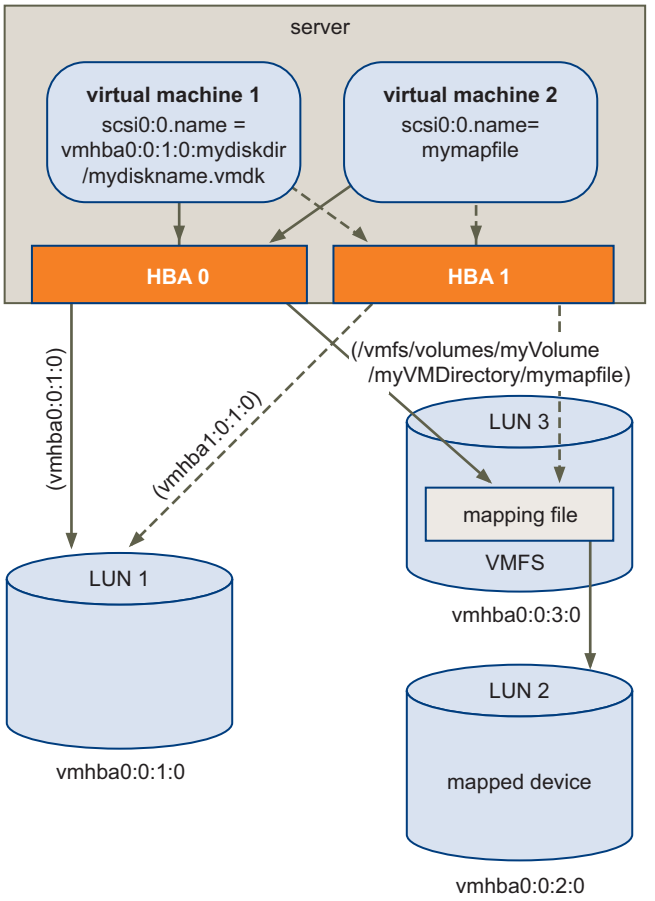

VMFS uniquely identifies all mapped LUNs, and the identification is stored in its internal data structures. Any change in the SCSI path, such as a Fibre Channel switch failure or the addition of a new host bus adapter, can change the vmhba device name, because the name includes the path designation (initiator, target, LUN). Dynamic name resolution compensates for these changes by adjusting the data structures to retarget LUNs to their new device names.

Raw Device Mapping with Virtual Machine Clusters

Use an RDM with virtual machine clusters that need to access the same raw LUN for failover scenarios. The setup is similar to that of a virtual machine cluster that accesses the same virtual disk file, but an RDM replaces the virtual disk file.

Figure 8-6. Access from Clustered Virtual Machines

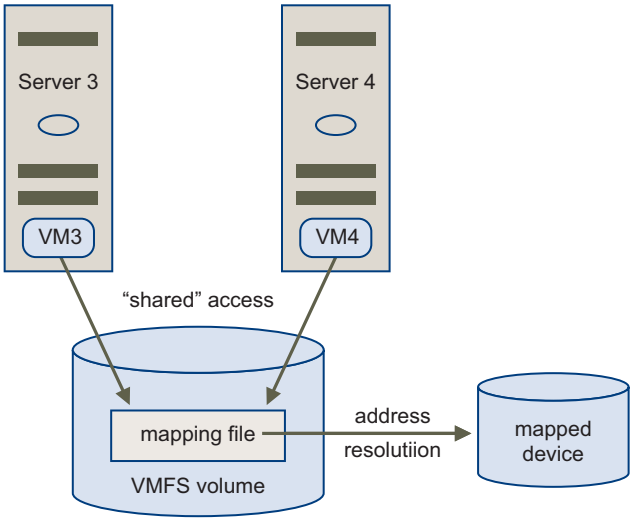

See the *Resource Management Guide*.

Comparing Raw Device Mapping to Other Means of SCSI Device Access

To help you choose among several available access modes for SCSI devices, Table 8-1 provides a quick comparison of features available with the different modes.

Table 8-1. Features Available with Virtual Disks and Raw Device Mappings

ESX Server 3 Features	Virtual Disk File	Virtual Mode RDM	Physical Mode RDM
SCSI Commands Passed Through	No	No	Yes REPORT LUNs is not passed through
VirtualCenter Support	Yes	Yes	Yes
Snapshots	Yes	Yes	No
Distributed Locking	Yes	Yes	Yes

Table 8-1. Features Available with Virtual Disks and Raw Device Mappings (Continued)

ESX Server 3 Features	Virtual Disk File	Virtual Mode RDM	Physical Mode RDM
Clustering	Cluster-in-a-box only	Cluster-in-a-box and cluster-across-boxes	N+1 (physical to virtual clustering only)
SCSI Target-Based Software	No	No	Yes

VMware recommends that you use virtual disk files for the cluster-in-a-box type of clustering. If you plan to reconfigure your cluster-in-a-box clusters as cluster-across-boxes clusters, use virtual mode RDMs for the cluster-in-a-box clusters. See the *Resource Management Guide*.

Managing Mapped LUNs

The tools available to manage mapped LUNs and their RDMs include the VI Client, the `vmkfstools` utility, and ordinary file system utilities used in the service console.

VMware Infrastructure Client

Using the VI Client, you can map a SAN LUN to a datastore and manage paths to your mapped LUN.

Creating Virtual Machines with RDMs

When you give your virtual machine direct access to a raw SAN LUN, you create a mapping file (RDM) that resides on a VMFS datastore and points to the LUN. Although the mapping file has the same `.vmdk` extension as a regular virtual disk file, the RDM file contains only mapping information. The actual virtual disk data is stored directly on the LUN.

You can create the RDM as an initial disk for a new virtual machine or add it to an existing virtual machine. When creating the RDM, you specify the LUN to be mapped and the datastore on which to put the RDM.

To create a virtual machine with an RDM

1 Follow all steps required to create a custom virtual machine.

See *Basic System Administration*.

2 In the Select a Disk page, select **Raw Device Mapping**, and click **Next**.

3 From the list of SAN disks or LUNs, select a raw LUN for your virtual machine to access directly.

For more information on configuring SAN storage, see the *Fibre Channel SAN Configuration Guide* or the *iSCSI SAN Configuration Guide*.

4 Select a datastore for the RDM mapping file.

You can place the RDM file on the same datastore where your virtual machine configuration file resides, or select a different datastore.

NOTE To use VMotion for virtual machines with enabled NPIV, make sure that the RDM files of the virtual machines are located on the same datastore. You cannot perform Storage VMotion, or VMotion between datastores, when NPIV is enabled.

---

5 Select a compatibility mode:

- **Physical compatibility** mode allows the guest operating system to access the hardware directly. Physical compatibility is useful if you are using SAN-aware applications on the virtual machine. However, a virtual machine with the physical compatibility RDM cannot be cloned, made into a template, or migrated if the migration involves copying the disk.
- **Virtual compatibility** mode allows the RDM to behave as if it were a virtual disk, so you can use such features as snapshotting, cloning, and so on.

6 Select a virtual device node.

7 If you select **Independent** mode, choose one of the following:

- **Persistent** – Changes are immediately and permanently written to the disk.
- **Nonpersistent** – Changes to the disk are discarded when you power off or revert to the snapshot.

8 Click **Next**.

9 In the Ready to Complete New Virtual Machine page, review your selections.

10 Click **Finish** to complete your virtual machine.

You can also add an RDM to an existing virtual machine.

To add an RDM to a virtual machine

1 From the VI Client, click **Inventory** in the navigation bar, and expand the inventory as needed.

2 Select the virtual machine from the inventory panel.

3 On the **Summary** tab, click **Edit Settings**.

4 Click **Add**.

The Add Hardware Wizard opens.

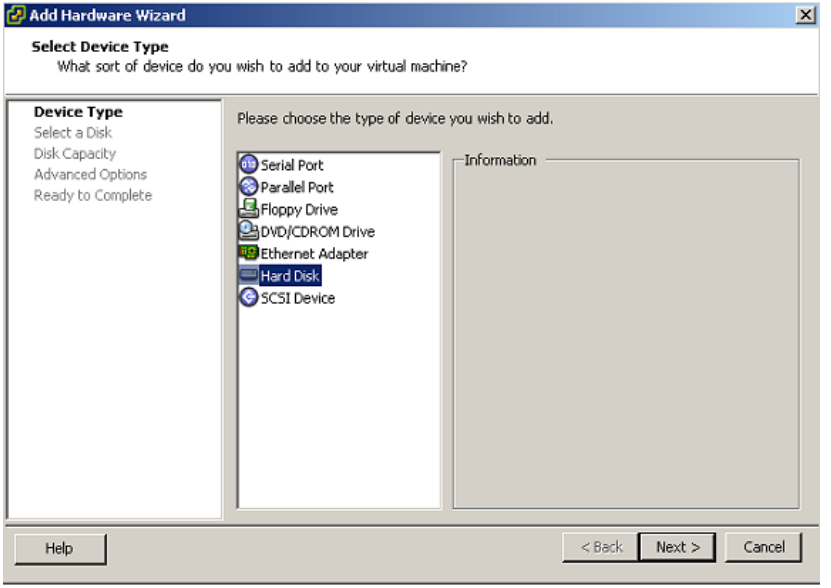

5 Choose **Hard Disk** as the type of device to add, and click **Next**.

6 Select **Raw Device Mapping**, and click **Next**.

7 Go to Step 3 of the preceding procedure.

Managing Paths for a Mapped Raw LUN

Use the Manage Paths dialog box to manage paths for your mapping files and mapped raw LUNs.

To manage paths

1 Log in as administrator or as the owner of the virtual machine to which the mapped disk belongs.

2 Select the virtual machine from the inventory panel.

3 On the **Summary** tab click the **Edit Settings** link.

The Virtual Machine Properties dialog box opens.

4 On the **Hardware** tab, select **Hard Disk,** then click **Manage Paths**.

5 Use the Manage Paths dialog box to enable or disable your paths, set multipathing policy, and specify the preferred path.

Follow these procedures:

- "To set the multipathing policy" on page 143
- "To set the preferred path" on page 144
- "To disable a path" on page 144

The vmkfstools Utility

The `vmkfstools` command-line utility can be used in the service console to perform many of the same operations available through the VI Client. Typical operations applicable to RDMs are the commands to create a mapping file, to query mapping information such as the name and identification of the mapped device, and to import or export a virtual disk.

See "Using vmkfstools" on page 283.

File System Operations

Most common file system operations done in the service console can be applied to RDMs.

Table 8-2. Commands Used in Service Console

Command	Description
`ls –l`	Shows the file name and permissions of the mapping file, while showing the length of the mapped device.
`du`	Shows the space used by the mapped device, rather than the mapping file.
`mv`	Renames the mapping file, but does not affect the mapped device.
`cp`	Copies the contents of a mapped device. You cannot use it to copy a virtual disk file to a mapped device. Instead, use the `vmkfstools` command.
`dd`	Copies data into or out of a mapped device. However, VMware recommends that you use `vmkfstools` import and export commands to copy data.

Security

Security for ESX Server 3 Systems

9

ESX Server is developed with a focus on strong security. This section provides you with an overview of how VMware ensures security in the ESX Server environment, addressing system architecture from a security standpoint and giving you a list of additional security resources.

This chapter contains the following sections:

- "ESX Server 3 Architecture and Security Features" on page 164
- "Security Resources and Information" on page 175

ESX Server 3 Architecture and Security Features

From a security perspective, VMware ESX Server 3 consists of four major components: the virtualization layer, the virtual machines, the service console, and the virtual networking layer. Figure 9-1 provides an overview of these components.

Figure 9-1. ESX Server 3 Architecture

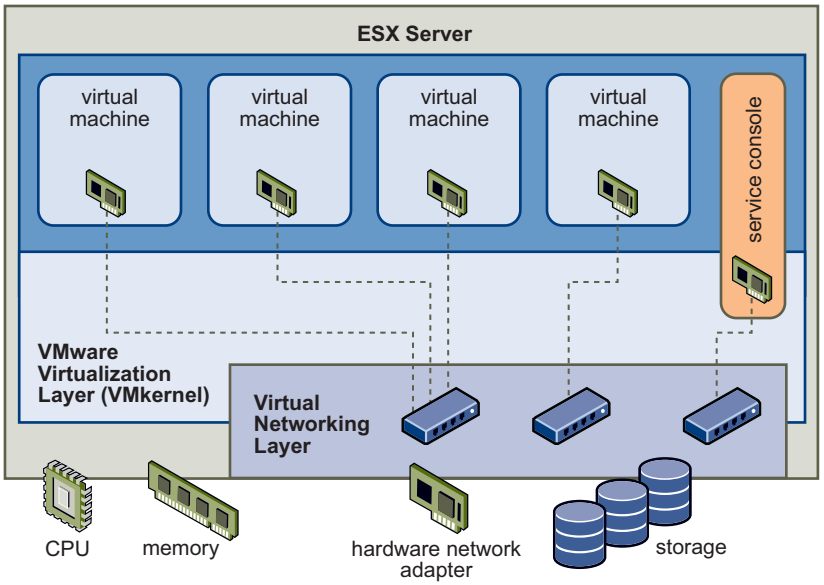

Each of these components and this overall architecture are designed to ensure security of the ESX Server 3 system as a whole.

Security and the Virtualization Layer

The virtualization layer, or VMkernel, is a kernel designed by VMware to run virtual machines. It controls the hardware that ESX Server hosts use and schedules the allocation of hardware resources among the virtual machines. Because the VMkernel is fully dedicated to supporting virtual machines and is not used for other purposes, the interface to the VMkernel is strictly limited to the API required to manage virtual machines.

Security and Virtual Machines

Virtual machines are the containers in which applications and guest operating systems run. All VMware virtual machines are isolated from one another. Virtual machines are designed to contain all users within the guest OS, regardless of the privileges they have. Even Administrators are isolated from other virtual machines in the same way that they are isolated from other physical machines.

This isolation enables multiple virtual machines to run securely while sharing hardware and ensures their ability to access hardware and their uninterrupted performance. For example, if a guest operating system running in a virtual machine crashes, other virtual machines on the same ESX Server host continue to run. The guest operating system crash has no effect on:

- The ability of users to access the other virtual machines
- The ability of the operational virtual machines to access the resources they need
- The performance of the other virtual machines

Each virtual machine is isolated from other virtual machines running on the same hardware. While virtual machines share physical resources such as CPU, memory, and I/O devices, a guest operating system on an individual virtual machine cannot detect any device other than the virtual devices made available to it, as shown in Figure 9-2.

Figure 9-2. Virtual Machine Isolation

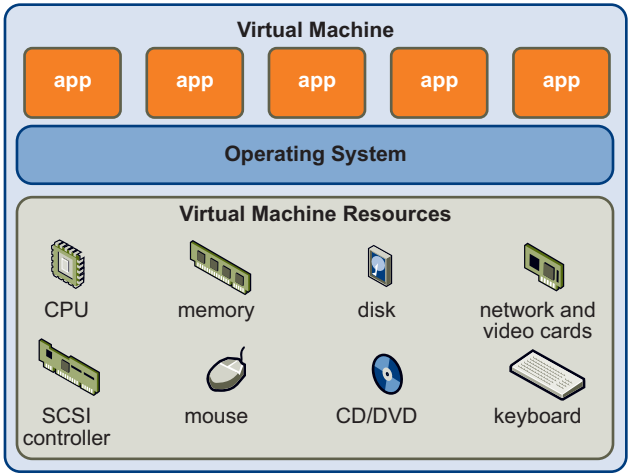

Because the VMkernel mediates the physical resources and all physical hardware access takes place through the VMkernel, virtual machines cannot circumvent this level of isolation.

Just as a physical machine can communicate with other machines in a network only through a network card, a virtual machine can communicate with other virtual machines running in the same ESX Server host only through a virtual switch. Further, a virtual machine communicates with the physical network, including virtual machines on other ESX Server hosts, only through a physical network adapter, as shown in Figure 9-3.

Figure 9-3. Virtual Networking Through Virtual Switches

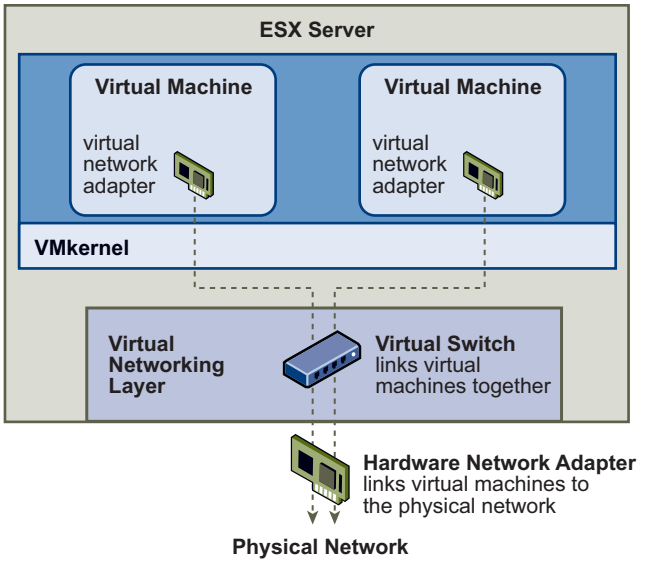

In considering virtual machine isolation in a network context, you can apply these rules:

- If a virtual machine does not share a virtual switch with any other virtual machine, it is completely isolated from virtual networks within the host.
- If no physical network adapter is configured for a virtual machine, the virtual machine is completely isolated from any physical networks.
- If you use the same safeguards (firewalls, antivirus software, and so forth) to protect a virtual machine from the network as you would for a physical machine, the virtual machine is as secure as the physical machine would be.

You can further protect virtual machines by setting up resource reservations and limits on the ESX Server host. For example, through the detailed resource controls available in ESX Server, you can configure a virtual machine so that it always receives at least ten percent of the ESX Server host's CPU resources, but never more than twenty percent.

Resource reservations and limits protect virtual machines from performance degradation if another virtual machine tries to consume too many resources on shared hardware. For example, if one of the virtual machines on an ESX Server host is incapacitated by a denial-of-service (DoS) or distributed denial-of-service (DDoS) attack, a resource limit on that machine prevents the attack from taking up so many hardware resources that the other virtual machines are also affected. Similarly, a resource reservation on each of the virtual machines ensures that, in the event of high resource demands by the virtual machine targeted by the DoS attack, all the other virtual machines still have enough resources to operate.

By default, ESX Server imposes a form of resource reservation by applying a distribution algorithm that divides the available host resources equally among the virtual machines while keeping a certain percentage of resources for use by other system components such as the service console. This default behavior provides a degree of natural protection from DoS and DDoS attacks. You set specific resource reservations and limits on an individual basis to customize the default behavior so that the distribution isn't equal across the virtual machine configuration. For a discussion of how to manage resource allocation for virtual machines, see the *Resource Management Guide*.

Security and the Service Console

The ESX Server 3 service console is a limited distribution of Linux based on Red Hat Enterprise Linux 3, Update 8 (RHEL 3 U8). The service console provides an execution environment to monitor and administer the entire ESX Server 3 host.

If the service console is compromised in certain ways, the virtual machines it interacts with might also be compromised. To minimize the risk of an attack through the service console, VMware protects the service console with a firewall. For information about this firewall, see "Service Console Firewall Configuration" on page 237.

In addition to implementing the service console firewall, the following are some of the other ways VMware mitigates risks to the service console:

- ESX Server 3 runs only services essential to managing its functions, and the distribution is limited to the features required to run ESX Server 3.
- By default, ESX Server 3 is installed with a high-security setting, which means that all outbound ports are closed and the only inbound ports that are open are those required for interactions with clients such as the VMware Virtual Infrastructure Client. VMware recommends that you keep this security setting unless the service console is connected to a trusted network.

- By default, all ports not specifically required for management access to the service console are closed. You must specifically open ports if you need additional services.
- All communications from clients are encrypted through SSL by default. The SSL connection uses 256-bit AES block encryption and 1024-bit RSA key encryption.
- The Tomcat Web service, used internally by ESX Server 3 to support access to the service console by Web clients like Virtual Infrastructure Web Access, has been modified to run only those functions required for administration and monitoring by a Web client. As a result, ESX Server 3 is not vulnerable to the Tomcat security issues reported in broader usage.
- VMware monitors all security alerts that could affect service console security and, if needed, issues a security patch, as it would for any other security vulnerability that could affect ESX Server 3 hosts. VMware provides security patches for RHEL 3 U6 and later as they become available.
- Insecure services such as FTP and Telnet are not installed and the ports for these services are closed by default. Because more secure services such as SSH and SFTP are easily available, always avoid using these insecure services in favor of their safer alternatives. If you must use insecure services and have implemented sufficient protection for the service console, you must explicitly open ports to support them.
- The number of applications that use a `setuid` or `setgid` flag is minimized, and you can disable any `setuid` or `setgid` application that is optional to ESX Server 3 operation. For information on required and optional `setuid` and `setgid` applications, see "setuid and setgid Applications" on page 250.

For details on these security measures and other service console security recommendations, see "Service Console Security" on page 235.

Although you can install and run certain types of programs designed for RHEL 3 U6 in the service console, this usage can have serious security consequences and is not supported unless VMware explicitly states that it is. If a security vulnerability is discovered in a supported configuration, VMware proactively notifies all customers with valid support and subscription contracts and provides all necessary patches.

NOTE Some security advisories issued by Red Hat do not apply to the ESX Server 3 environment. VMware does not provide notification or patches in these instances.

To learn more about VMware's policies on security patches for supported programs, as well as its policies on unsupported software, see "Security Resources and Information" on page 175.

Security and the Virtual Networking Layer

The virtual networking layer consists of the virtual network devices through which virtual machines and the service console interface with the rest of the network. ESX Server relies on the virtual networking layer to support communications between virtual machines and their users. In addition, ESX Server hosts use the virtual networking layer to communicate with iSCSI SANs, NAS storage, and so forth. The virtual networking layer includes virtual network adapters and the virtual switches.

The methods you use to secure a virtual machine network depend on which guest operating system is installed, whether the virtual machines operate in a trusted environment, and a variety of other factors. Virtual switches provide a substantial degree of protection when used with other common security practices, such as installing firewalls. ESX Server also supports IEEE 802.1q VLANs, which you can use to further protect the virtual machine network, service console, or storage configuration. VLANs let you segment a physical network so that two machines on the same physical network cannot send packets to or receive packets from each other unless they are on the same VLAN.

You can get a sense of how to use virtual switches to implement security tools like DMZs and configure virtual machines on different networks within the same ESX Server host by reviewing the following examples.

NOTE For a specific discussion of how virtual switches and VLANs help safeguard the virtual machine network and a discussion of other security recommendations for virtual machine networks, see "Securing Virtual Machines with VLANs" on page 195.

Example: Creating a Network DMZ On a Single ESX Server Host

One example of how to use ESX Server isolation and virtual networking features to configure a secure environment is the creation of a network demilitarized zone (DMZ) on a single ESX Server host, as shown in Figure 9-4.

Figure 9-4. DMZ Configured On a Single ESX Server Host

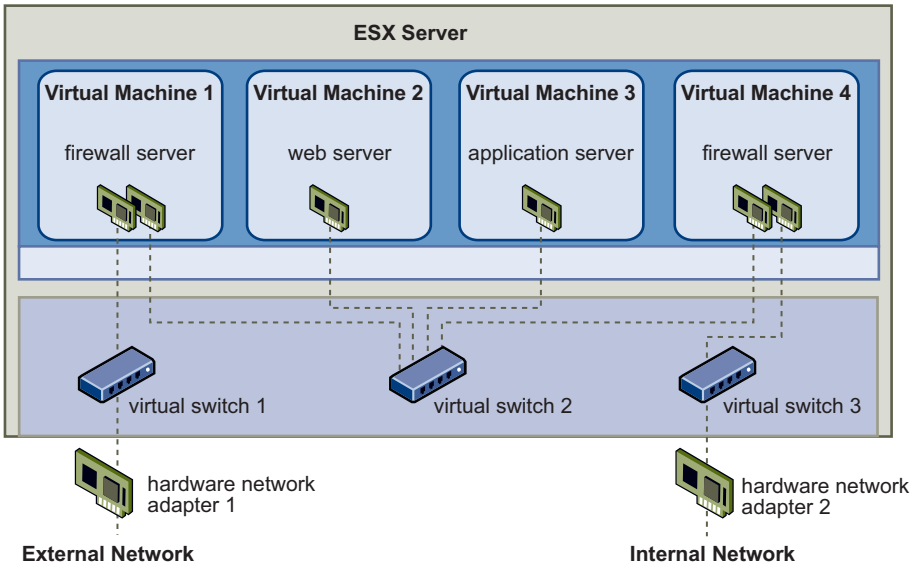

This configuration includes four virtual machines configured to create a virtual DMZ on Virtual Switch 2. Virtual Machine 1 and Virtual Machine 4 run firewalls and are connected to virtual adapters through virtual switches. Both of these virtual machines are multihomed. Of the remaining two virtual machines, Virtual Machine 2 runs a Web server and Virtual Machine 3 runs as an application server. Both of these virtual machines are single homed.

The Web server and application server occupy the DMZ between the two firewalls. The conduit between these elements is Virtual Switch 2, which connects the firewalls with the servers. This switch has no direct connection with any elements outside the DMZ and is isolated from external traffic by the two firewalls.

From an operational viewpoint, external traffic from the Internet enters Virtual Machine 1 through Hardware Network Adapter 1 (routed by Virtual Switch 1) and is verified by the firewall installed on this machine. If the firewall authorizes the traffic, it is routed to the virtual switch in the DMZ, Virtual Switch 2. Because the Web server and application server are also connected to this switch, they can serve external requests.

Virtual Switch 2 is also connected to Virtual Machine 4. This virtual machine provides a firewall between the DMZ and the internal corporate network. This firewall filters packets from the Web server and application server. If a packet is verified, it is routed to Hardware Network Adapter 2 through Virtual Switch 3. Hardware Network Adapter 2 is connected to the internal corporate network.

When creating a DMZ on a single ESX Server, you can use fairly lightweight firewalls. Although a virtual machine in this configuration cannot exert direct control over another virtual machine or access its memory, all the virtual machines are still connected through a virtual network, and this network could be used for virus propagation or targeted for other types of attacks. The virtual machines in the DMZ as secure as separate physical machines connected to the same network.

Example: Creating Multiple Networks Within a Single ESX Server Host

The ESX Server system is designed so that you can connect some groups of virtual machines to the internal network, others to the external network, and still others to both—all on the same ESX Server host. This capability is an outgrowth of basic virtual machine isolation coupled with a well-planned use of virtual networking features, as shown in Figure 9-5.

Figure 9-5. External Networks, Internal Networks, and a DMZ Configured On a Single ESX Server Host

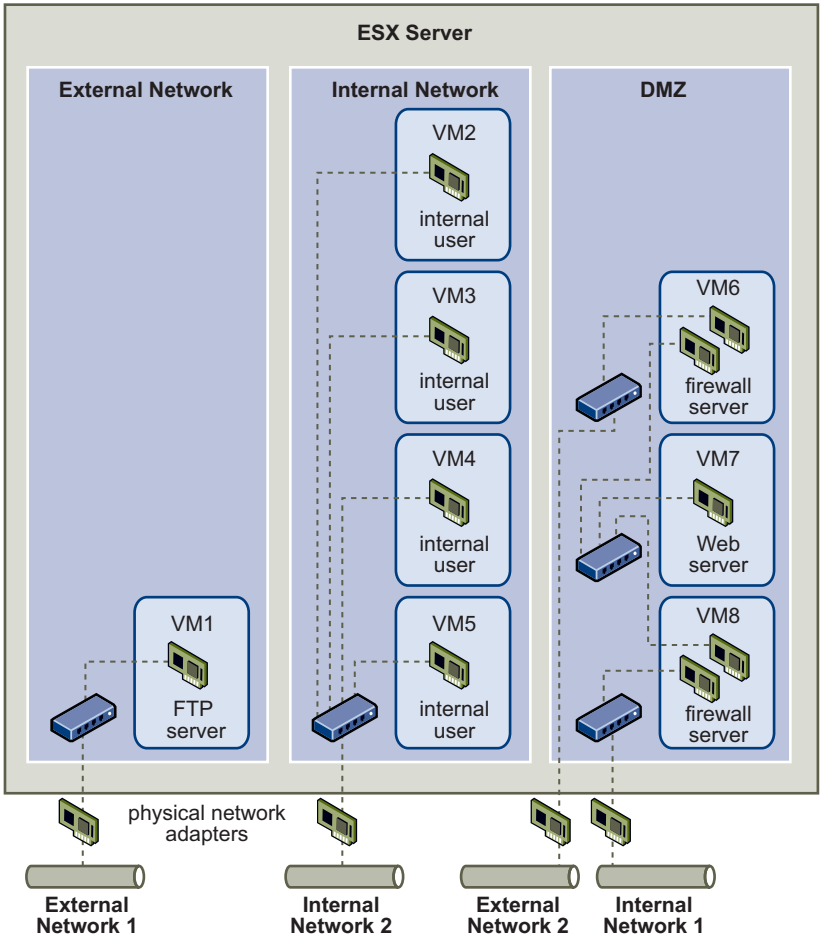

Here, the system administrator configured an ESX Server host into three distinct virtual machine zones, each serving a unique function:

- **FTP server** – Virtual Machine 1 is configured with FTP software and acts as a holding area for data sent to and from outside resources such as forms and collateral localized by a vendor.

  This virtual machine is associated with an external network only. It has its own virtual switch and physical network adapter that connect it to External Network 1. This network is dedicated to servers that the company uses to receive data from outside sources. For example, the company uses External Network 1 to receive FTP traffic from vendors and allow vendors access to data stored on externally available servers though FTP. In addition to servicing Virtual Machine 1, External Network 1 services FTP servers configured on different ESX Server hosts throughout the site.

  Because Virtual Machine 1 doesn't share a virtual switch or physical network adapter with any virtual machines in the host, the other resident virtual machines cannot transmit packets to or receive packets from the Virtual Machine 1 network. This restriction prevents sniffing attacks, which require sending network traffic to the victim. More importantly, an attacker cannot use the natural vulnerability of FTP to access any of the host's other virtual machines.

- **Internal virtual machines** – Virtual Machines 2 through 5 are reserved for internal use. These virtual machines process and store company-private data such as medical records, legal settlements, and fraud investigations. As a result, the system administrators must ensure the highest level of protection for these virtual machines.

  These virtual machines connect to Internal Network 2 through their own virtual switch and network adapter. Internal Network 2 is reserved for internal use by personnel such as claims processors, in-house lawyers, or adjustors.

  Virtual Machines 2 through 5 can communicate with one another through the virtual switch and with internal virtual machines elsewhere on Internal Network 2 through the physical network adapter. They cannot communicate with externally-facing machines. As with the FTP server, these virtual machines cannot send packets to or receive packets from the other virtual machines' networks. Similarly, the host's other virtual machines cannot send packets to or receive packets from Virtual Machines 2 through 5.

- **DMZ** – Virtual Machines 6 through 8 are configured as a DMZ that the marketing group uses to publish the company's external Web site.

  This group of virtual machines is associated with External Network 2 and Internal Network 1. The company uses External Network 2 to support the Web servers that use the marketing and financial department to host the corporate Web site and other Web facilities that it hosts to outside users. Internal Network 1 is the conduit that the marketing department uses to publish web pages to the corporate Web site, post downloads, and maintain services like user forums.

  Because these networks are separate from External Network 1 and Internal Network 2 and the virtual machines have no shared points of contact (switches or adapters), there is no risk of attack to or from the FTP server or the internal virtual machine group.

  For an example of how to configure a DMZ with virtual machines, see "Example: Creating a Network DMZ On a Single ESX Server Host" on page 170.

By capitalizing on virtual machine isolation, correctly configuring virtual switches, and maintaining network separation, the system administrator can house all three virtual machine zones in the same ESX Server host and be confident that there will be no data or resource breaches.

The company enforces isolation among the virtual machine groups by using multiple internal and external networks and making sure that the virtual switches and physical network adapters for each group are completely separate from those of other groups.

Because none of the virtual switches straddle virtual machine zones, the system administrator succeeds in eliminating the risk of packet leakage from one zone to another. A virtual switch, by design, cannot leak packets directly to another virtual switch. The only way for packets to travel from one virtual switch to another is if:

- The virtual switches are connected to the same physical LAN.
- The virtual switches connect to a common virtual machine, which could then be used to transmit packets.

Neither of these conditions occur in the sample configuration. If system administrators want to verify that no common virtual switch paths exist, they can check for possible shared points of contact by reviewing the network switch layout in the VI Client or VI Web Access. For information on the virtual switch layout, see "Virtual Switches" on page 21.

To safeguard the virtual machines' resources, the system administrator lowers the risk of DoS and DDoS attacks by configuring a resource reservation and limit for each virtual machine. The system administrator further protects the ESX Server host and virtual machines by installing software firewalls at the front and back ends of the DMZ, ensuring that the ESX Server host is behind a physical firewall, and configuring the service console and networked storage resources so that each has its own virtual switch.

Security Resources and Information

You can find additional information on security topics through the following resources.

Table 9-1. VMware Security Resources on the Web

Topic	Resource
VMware security policy, up-to-date security alerts, security downloads, and focus discussions of security topics	www.vmware.com/security/ communities.vmware.com/community/vmtn/general/security
Corporate security response policy	http://www.vmware.com/support/policies/security_response.html VMware is committed to helping you maintain a secure environment. To reassure you that any security issues will be corrected in a timely fashion, the VMware Security Response Policy states our commitment to resolve possible vulnerabilities in our products.
Third-party software support policy	http://www.vmware.com/support/policies VMware supports a variety of storage systems, software agents such as backup agents, system management agents, and so forth. You can find lists of agents, tools, and other software that supports ESX Server 3 by searching http://www.vmware.com/vmtn/resources for ESX Server 3 compatibility guides. The industry offers more products and configurations than VMware can test. If VMware does not list a product or configuration in a compatibility guide, Technical Support will attempt to help you with any problems, but cannot guarantee that the product or configuration can be used. Always evaluate any security risks for unsupported products or configurations carefully.

Securing an ESX Server 3 Configuration 10

This chapter describes measures you can take to promote a secure environment for your ESX Server 3 hosts, virtual machines, and iSCSI SANs. The discussion focuses on network configuration planning from a security perspective and the steps you can take to protect the components in your configuration from attack.

This chapter discusses the following topics:

- "Securing the Network with Firewalls" on page 177
- "Securing Virtual Machines with VLANs" on page 195
- "Securing iSCSI Storage" on page 204

Securing the Network with Firewalls

Security administrators use firewalls to safeguard the network or selected components in the network from intrusion. Firewalls control access to devices within their perimeter by closing all communication pathways except for those that the administrator explicitly or implicitly designates as authorized, thus preventing unauthorized use of the devices. The pathways, or ports, that administrators open in the firewall allow traffic between devices on different sides of the firewall.

In a virtual machine environment, you can plan your layout for firewalls between:

- Physical machines such as VirtualCenter Management Server hosts and ESX Server 3 hosts.
- One virtual machine and another—for example, between a virtual machine acting as an external Web server and a virtual machine connected to your company's internal network.
- A physical machine and a virtual machine such as when you place a firewall between a physical network adapter card and a virtual machine.

How you use firewalls in an ESX Server 3 configuration is based on how you plan to use the network and how secure any given component needs to be. For example, if you create a virtual network where each virtual machine is dedicated to running a different benchmark test suite for the same department, the risk of unwanted access from one virtual machine to the next is minimal. Hence, you have little need to set up the configuration so that firewalls are present between the virtual machines. However, to prevent interruption of a test run from an outside host, you might set up the configuration so that a firewall is present at the entry point of the virtual network to protect the entire set of virtual machines.

This section shows firewall placement for configurations with and without VirtualCenter. It also provides information on the firewall ports required for ESX Server 3 systems.

- "Firewalls for Configurations with a VirtualCenter Server" on page 179
- "Firewalls for Configurations Without a VirtualCenter Server" on page 182
- "TCP and UDP Ports for Management Access" on page 183
- "Connecting to VirtualCenter Server Through a Firewall" on page 185
- "Connecting to the Virtual Machine Console Through a Firewall" on page 186
- "Connecting ESX Server 3 Hosts Through Firewalls" on page 188
- "Opening Firewall Ports for Supported Services and Management Agents" on page 188

For information on the service console firewall, see "Service Console Firewall Configuration" on page 237. To configure the firewall and port settings during installation, see the *Setup Guide*.

Firewalls for Configurations with a VirtualCenter Server

If you use a VirtualCenter Server, you can install firewalls at any of the locations shown in Figure 10-1.

NOTE Depending on your configuration, you might not need all the firewalls in the illustration, or you might need firewalls in locations not shown.

Figure 10-1. Sample Virtual Infrastructure Network Configuration and Traffic Flow

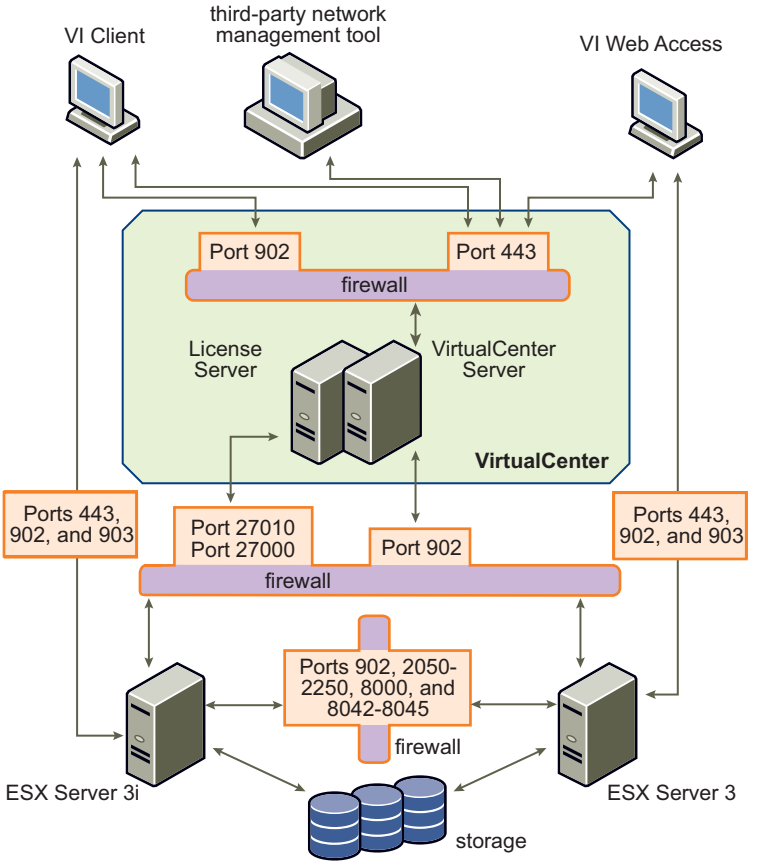

Networks configured with a VirtualCenter Server can receive communications through several types of clients: the VI Client, VI Web Access, or third-party network management clients that use the SDK to interface with the host. During normal operation, VirtualCenter listens for data from its managed hosts and clients on designated ports. VirtualCenter also assumes that its managed hosts listen for data from VirtualCenter on designated ports. If a firewall is present between any of these elements, you must ensure that the firewall has open ports to support data transfer.

If you access ESX Server 3 hosts through a VirtualCenter Server, you typically protect the VirtualCenter Server by using a firewall. This firewall provides basic protection for your network. Whether this firewall lies between the clients and the VirtualCenter Server or the VirtualCenter Server and the clients are behind the firewall depends on your deployment. The main point is to ensure that a firewall is present at what you consider to be an entry point for the system as a whole.

You might also include firewalls at a variety of other access points in the network, depending on how you plan to use the network and how secure the various devices need to be. Select the locations for your firewalls based on the security risks that you've identified for your network configuration. The following is a list of firewall locations common to ESX Server 3 implementations. Many of the firewall locations in the list and Figure 10-1 are optional.

- Between your Web browser and VI Web Access HTTP and HTTPS proxy server.
- Between the VI Client, VI Web Access Client, or a third-party network-management client and the VirtualCenter Server.
- If your users access virtual machines through the VI Client, between the VI Client and the ESX Server 3 host. This connection is in addition to the connection between the VI Client and the VirtualCenter Server, and it requires a different port.
- If your users access virtual machines through a Web browser, between the Web browser and the ESX Server 3 host. This connection is in addition to the connection between the VI Web Access Client and VirtualCenter Server, and it requires different ports.
- Between the license server and either the VirtualCenter Server or the ESX Server 3 host. Typically, in configurations that include a VirtualCenter Server, the license server runs on the same physical machine as does the VirtualCenter Server. In this case, the license server connects to the ESX Server 3 network through a firewall, running in parallel with the VirtualCenter Server but using different ports.

  In some configurations, you might use an external license server—for example, if your company wants to control all licenses through a single, dedicated appliance. Here, you would connect the license server to the VirtualCenter Server through a firewall between these two servers.

Regardless of how you set up the license-server connection, the ports you use for license traffic are the same. For information on licensing, see the *Setup Guide*.

- Between the VirtualCenter Server and the ESX Server 3 hosts.
- Between the ESX Server 3 hosts in your network. Although traffic between ESX Server 3 hosts is usually considered trusted, you can add firewalls between your ESX Server 3 hosts if you are concerned about security breaches from machine to machine.

  If you add firewalls between ESX Server 3 hosts and plan to migrate virtual machines between the servers, perform cloning, or use VMotion, you must also open ports in any firewall that divides the source host from the target hosts so that the source and targets can communicate.
- Between the ESX Server 3 hosts and network storage such as NFS or iSCSI storage. These ports are not specific to VMware, and you configure them according to the specifications for your network.

For information on the ports to open for these communications paths, see "TCP and UDP Ports for Management Access" on page 183.

Firewalls for Configurations Without a VirtualCenter Server

If you connect clients directly to your ESX Server 3 network instead of using a VirtualCenter Server, your firewall configuration is somewhat simpler. You might install firewalls at any of the locations shown in Figure 10-2.

NOTE Depending on your configuration, you might not need all the firewalls in the Figure 10-2, or you might need firewalls in locations not shown.

Figure 10-2. Firewall Configuration for ESX Server 3 Networks that a Client Manages Directly

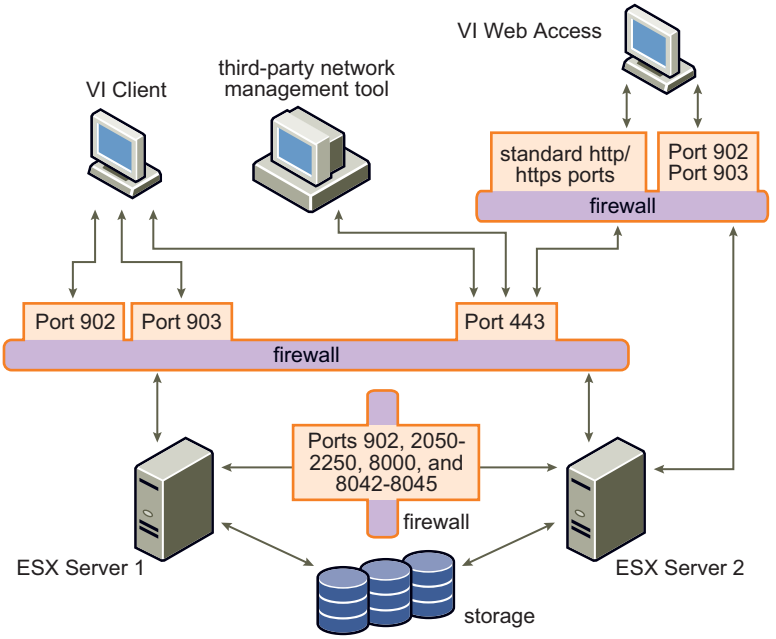

Networks configured without a VirtualCenter Server receive communications through the same types of clients as they do if a VirtualCenter Server were present: VI Client, VI Web Access Clients, or third-party network management clients. For the most part, the firewall needs are the same, but there are several key differences:

- Just as you would for configurations that include a VirtualCenter Server, be sure a firewall is present to protect your ESX Server 3 layer or, depending on your configuration, your clients and ESX Server 3 layer. This firewall provides basic protection for your network. The firewall ports you use are the same as those you use if a VirtualCenter Server is in place.
- Licensing in this type of configuration is part of the ESX Server 3 package that you install on each of the ESX Server 3 hosts. Because licensing is resident to the server, you do not need to install a separate license server. This eliminates the need for a firewall between the license server and the ESX Server 3 network.

NOTE In some situations, you might want to centralize your licenses. You can choose to maintain a separate license server or house the license server on one of the ESX Server 3 hosts in your network. With either of these approaches, you connect the license server to the ESX Server 3 network through a firewall by using the ports normally reserved for virtual machine licensing, much as you do if a VirtualCenter Server is present. Configurations that use a license server other than the one automatically installed on the ESX Server 3 host require additional setup. For information on licensing, see the *Setup Guide*.

TCP and UDP Ports for Management Access

This section lists predetermined TCP and UDP ports used for management access to your VirtualCenter Server, ESX Server 3 hosts, and other network components. To manage network components from outside a firewall, you might need to reconfigure the firewall to allow access on the appropriate ports.

NOTE The ports listed in Table 10-1 are connected through the service console interface unless otherwise indicated.

Table 10-1. TCP and UDP Ports

Port	Purpose	Traffic Type
80	HTTP access. The default non-secure TCP Web port typically used in conjunction with port 443 as a front end for access to ESX Server 3 networks from the Web, Port 80 redirects traffic to an HTTPS landing page (port 443) from which you launch the virtual machine console. Use port 80 for connection to VI Web Access from the Web. WS-Management uses port 80.	Incoming TCP
427	The CIM client uses the Service Location Protocol, version 2 (SLPv2) to find CIM servers.	Incoming and outgoing UDP
443	HTTPS access. The default SSL Web port. Use Port 443 for the following: ■ Connection to VI Web Access from the Web. ■ VI Web Access and third-party network management client connections to the VirtualCenter Server. ■ Direct VI Web Access and third-party network management clients access to ESX Server 3 hosts. ■ VI Client access to the VirtualCenter Server. ■ Direct VI Client access to ESX Server 3 hosts. ■ WS-Management. ■ VMware Update Manager. ■ VMware Converter.	Incoming TCP
902	Authentication traffic for the ESX Server 3. Use Port 902 for ESX Server 3 host access to other ESX Server 3 hosts for migration and provisioning.	Incoming TCP, outgoing UDP
903	Remote console traffic that user access to virtual machines generates on a specific ESX Server 3 host. Use Port 903 for the following: ■ VI Client access to virtual machine consoles. ■ VI Web Access Client access to virtual machine consoles.	Incoming TCP
2049	Transactions from your NFS storage devices. This port is used on the VMkernel interface rather than the service console interface.	Incoming and outgoing TCP
2050–2250	Traffic between ESX Server 3 hosts for VMware High Availability (HA) and EMC Autostart Manager.	Outgoing TCP, incoming and outgoing UDP
3260	Transactions from your iSCSI storage devices. This port is used on the VMkernel interface and the service console interface.	Outgoing TCP

Table 10-1. TCP and UDP Ports (Continued)

Port	Purpose	Traffic Type
5900-5906	RFB protocol which is used by management tools such as VNC.	Incoming and outgoing TCP
5988	CIM XML transactions over HTTPS. (SEE UPDATE)	Incoming and outgoing TCP
5989	CIM XML transactions over HTTP. (SEE UPDATE)	Incoming and outgoing TCP
8000	Incoming requests from VMotion. This port is used on the VMkernel interface rather than the service console interface.	Incoming and outgoing TCP
8042–8045	Traffic between ESX Server 3 hosts for HA and EMC Autostart Manager.	Outgoing TCP, incoming and outgoing UDP
27000	License transactions from ESX Server 3 to the license server (lmgrd.exe)	Incoming and outgoing TCP
27010	License transactions from ESX Server 3 to the license server (vmwarelm.exe).	Incoming and outgoing TCP

In addition to the TCP and UDP ports just discussed, you can configure other ports depending on your needs:

- You can use VI Client to open ports for installed management agents and supported services such as SSH, NFS, and so forth. For information on configuring additional ports for these services, see "Opening Firewall Ports for Supported Services and Management Agents" on page 188.
- You can open ports in the service console firewall for other services and agents required for your network by running command-line scripts. See "Service Console Firewall Configuration" on page 237.

Connecting to VirtualCenter Server Through a Firewall

As shown in Table 10-1, the port that VirtualCenter Server uses to listen for data transfer from its clients is 443. If you have a firewall between your VirtualCenter Server and its clients, you must configure a connection through which the VirtualCenter Server can receive data from the clients.

To enable the VirtualCenter Server to receive data from a Client, open port 443. Contact the firewall system administrator for additional information on configuring ports in a firewall.

If you are using the VI Client and do not want to use port 443 as the port for the VI Client-to-VirtualCenter Server communication, you can switch to another port by changing the VirtualCenter settings in the VI Client. To learn how to change these settings, see the *Basic System Administration Guide*.

Connecting to the Virtual Machine Console Through a Firewall

Whether you connect your client to ESX Server 3 hosts through a VirtualCenter Server or use a direct connection to the ESX Server 3 host, certain ports are required for user and administrator communication with virtual machine consoles. These ports support different client functions, interface with different layers on ESX Server 3, and use different authentication protocols. They are:

- **Port 902** – The VirtualCenter Server uses this port to send data to the VirtualCenter managed hosts. Port 902 is the port that the VirtualCenter Server assumes is available when sending data to the ESX Server 3 host. VMware doesn't support configuring a different port for this connection.

  Port 902 connects the VirtualCenter Server to the ESX Server 3 host through the VMware Authorization Daemon (`vmware-authd`). This daemon then multiplexes port 902 data to the appropriate recipient for processing.

- **Port 443** – The VI Client, VI Web Access Client, and SDK use this port to send data to the VirtualCenter managed hosts. Also, the VI Client, VI Web Access Client, and SDK, when connected directly to an ESX Server 3 host, use this port to support any management functions related to the server and its virtual machines. Port 443 is the port that clients assume is available when sending data to the ESX Server 3 host. VMware does not support configuring a different port for these connections.

  Port 443 connects clients to the ESX Server 3 host through the Tomcat Web service or the SDK. The `vmware-hostd` multiplexes port 443 data to the appropriate recipient for processing.

- **Port 903** – The VI Client and VI Web Access use this port to provide a connection for guest operating system mouse-keyboard-screen (MKS) activities on virtual machines. It is through this port that users interact with the guest operating systems and applications of the virtual machine. Port 903 is the port that the VI Client and VI Web Access assume is available when interacting with virtual machines. VMware does not support configuring a different port for this function.

  Port 903 connects the VI Client to a specified virtual machine configured on the ESX Server 3 host.

Figure 10-3 shows the relationships between VI Client functions, ports, and ESX Server 3 processes. The VI Web Access Client uses the same basic mapping for its interactions with the ESX Server 3 host.

Figure 10-3. Port Usage for VI Client Communications with ESX Server 3

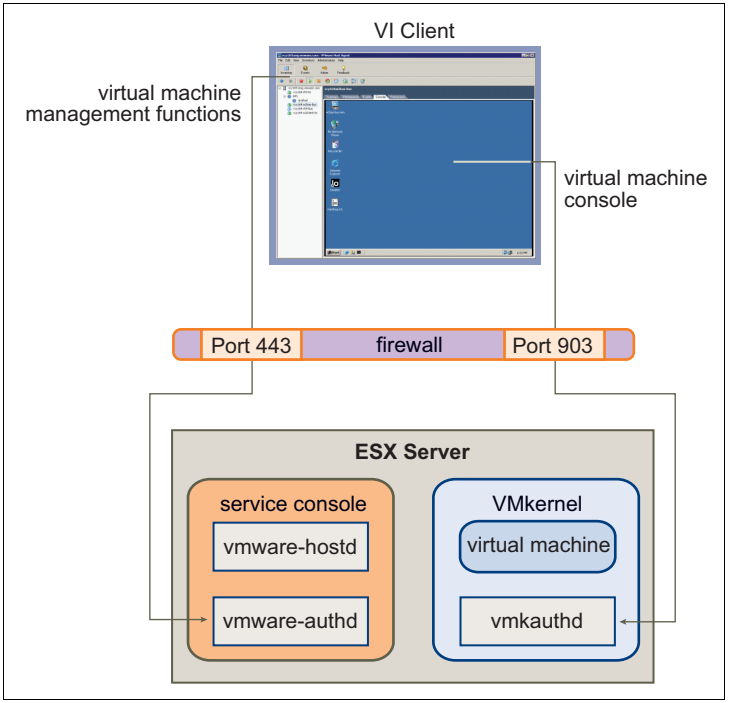

If you have a firewall between your VirtualCenter Server and VirtualCenter managed host, open Ports 443 and 903 in the firewall to allow data transfer to:

- ESX Server 3 hosts from the VirtualCenter Server.
- ESX Server 3 hosts directly from the VI Client and VI Web Access.

For additional information on configuring the ports, see the firewall system administrator.

Connecting ESX Server 3 Hosts Through Firewalls

If you have a firewall between two ESX Server 3 hosts and you want to allow transactions between the hosts or use VirtualCenter to perform any source or target activities, such as VMware High Availability (HA) traffic, migration, cloning, or VMotion, you must configure a connection through which the managed hosts can receive data. To do so, you open ports in the following ranges:

- 443 (server-to-server migration and provisioning traffic)
- 2050–2250 (for HA traffic)
- 8000 (for VMotion)
- 8042–8045 (for HA traffic)

For additional information on configuring the ports, see the firewall system administrator. For more detailed information on the directionality and protocol for these ports, see "TCP and UDP Ports for Management Access" on page 183.

Opening Firewall Ports for Supported Services and Management Agents

Use VI Client to configure the service console firewall to accept commonly supported services and installed management agents. When you configure the ESX Server 3 host security profile in VirtualCenter, you add or remove these services or agents, automatically opening or closing predetermined ports in the firewall to allow communication with the service or agent. The following is a list of the services and agents you can add or remove:

- NIS client
- NFS client (insecure service)
- SMB client (insecure service)
- FTP client (insecure service)
- SSH client
- Telnet client (insecure service)
- NTP client
- iSCSI software client
- SSH server
- Telnet server (insecure service)

- FTP server (insecure service)
- NFS server (insecure service)
- CIM HTTP server (insecure service)
- CIM HTTPS server
- SNMP server
- Syslog client
- Other supported management agents you install

NOTE This list can change, so you might find that the VI Client provides services and agents not mentioned in the list. Also, not all services on the list are installed by default. You might need to perform additional tasks to configure and enable these services.

If you are installing a device, service, or agent not on this list, open ports in the service console firewall from a command line. See "Service Console Firewall Configuration" on page 237.

To allow access to ESX Server 3 for a service or management agent

1 Log in to the VI Client and select the server from the inventory panel.

2 Click the **Configuration** tab and click **Security Profile**.

The VI Client displays a list of currently active incoming and outgoing connections with the corresponding firewall ports.

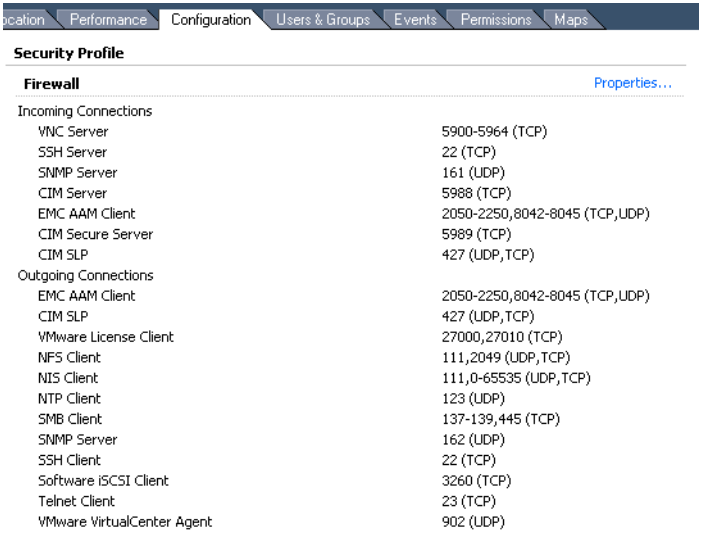

3 Click **Properties** to open the **Firewall Properties** dialog box.

The **Firewall Properties** dialog box lists all the services and management agents you can configure for the host.

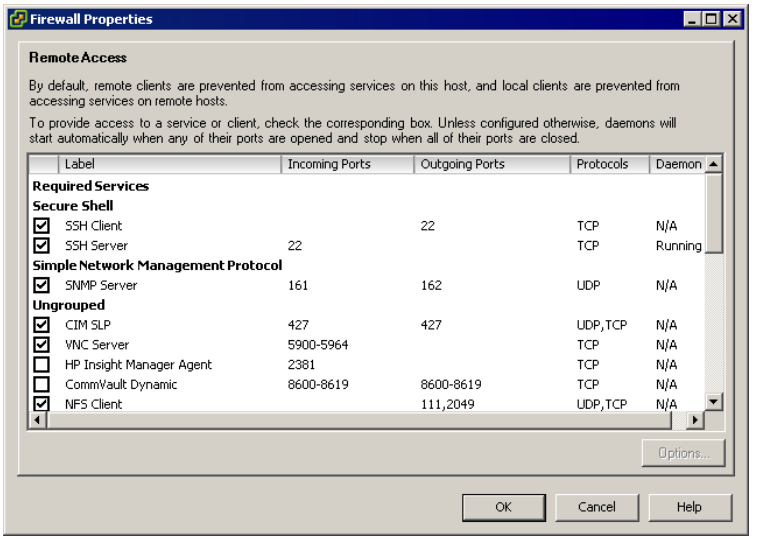

4 Select the services and agents to enable.

The **Incoming Ports** and **Outgoing Ports** columns indicate the port or ports that the VI Client opens for the service, the **Protocol** column indicates the protocol the service uses, and the **Daemon** column indicates the status of daemons associated with the service.

5 Click **OK**.

Automating Service Behavior Based on Firewall Settings

ESX Server 3 can automate whether or not services start based on the status of firewall ports. Such automation helps ensure that services start if the environment is configured to enable their function. For example, starting a network service only if some ports are open can help avoid the situation where services are started, but are unable to complete the communications required to complete their intended purpose.

For example, having accurate information about the current time is a requirement for some protocols, such as Kerberos. The NTP service is a way of getting accurate time information, but this service only works when required ports are opened in the firewall. Therefore, the service cannot achieve its goal if all ports are closed. The NTP services

provide an option to configure the conditions when the service starts or stops. This configuration includes options that account for whether firewall ports are opened, and then start or stop the NTP service based on those conditions. Several possible configuration options exist, all of which are also applicable to the SSH server.

NOTE The settings described in this section only apply to service settings configured through the VI Client or applications created with the VMware Infrastructure SDK. Configurations made through other means, such as the esxcfg-firewall utility or configuration files in `/etc/init.d/` are not affected by these settings.

- **Start automatically if any ports are open, and stop when all ports are closed** – The default setting for these services that VMware recommends. If any port is open, the client attempts to contact the network resources pertinent to the service in question. If some ports are open, but the port for a particular service is closed, the attempt fails, but there is little drawback to such a case, and if and when the applicable outgoing port is opened, the service begins completing its tasks.
- **Start and stop with host** – The service starts shortly after the host starts and closes shortly before the host shuts down. Much like **Start automatically if any ports are open, and stop when all ports are closed**, this option means that the service regularly attempts to complete its tasks, such as contacting the specified NTP server, in the case of NTP. If the port was closed, but is subsequently opened, the client begins completing its tasks shortly thereafter.
- **Start and stop manually** – The host preserves the user-determined service settings, regardless of whether ports are open or not. When a user starts the NTP service, that service is kept running as long as the host is powered on. If the service is started and the host is powered off, the service is stopped as part of the shut-down process, but as soon as the host is powered on, the service is started again, preserving the user-determined state.

To configure how service startup relates to firewall configuration

1 Log in to the VI Client and select the server from the inventory panel.

2 Click the **Configuration** tab and click **Security Profile**.

The VI Client displays a list of currently active incoming and outgoing connections with the corresponding firewall ports.

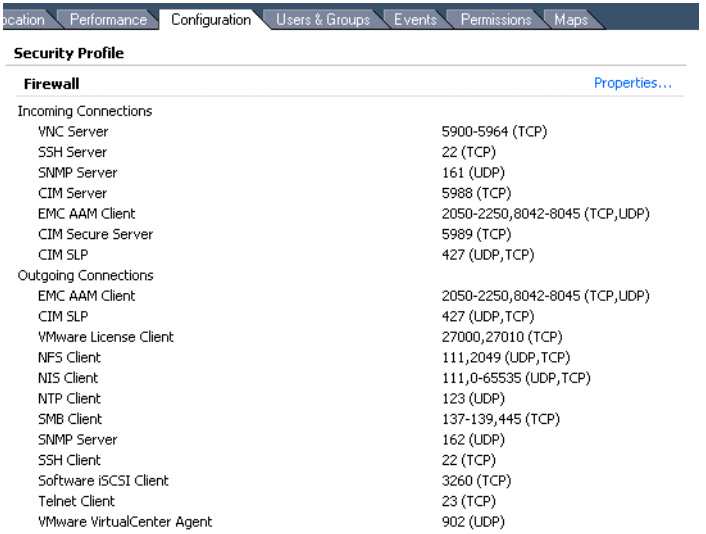

3 Click **Properties**.

The **Firewall Properties** dialog box lists all the services and management agents you can configure for the host.

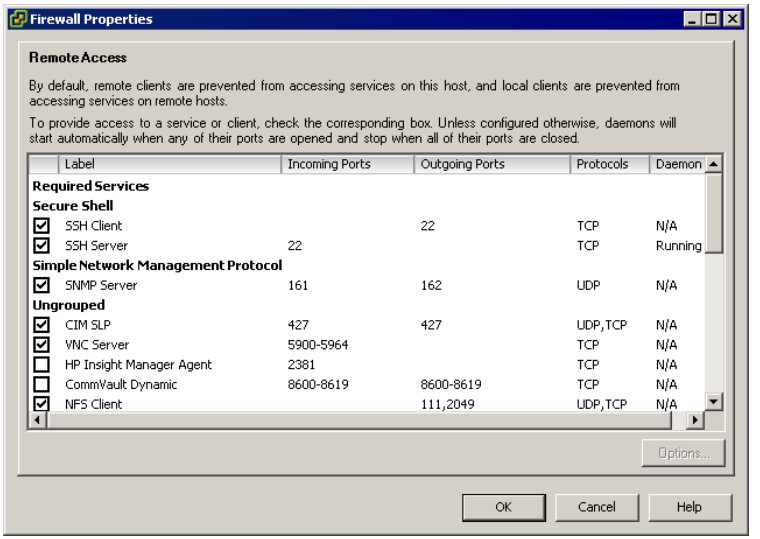

4 Select the service to configure, and click **Options**.

The **Startup Policy** determines when the service starts. This dialog box also provides information about the current state of the service and provides an interface for manually starting, stopping, or restarting the service.

5 Select an option from the **Startup Policy** choices.

6 Click **OK**.

Securing Virtual Machines with VLANs

The network can be one of the most vulnerable parts of any system. Just as the physical network requires protection, so does your virtual machine network. If your virtual machine network is connected to a physical network, it can be subject to breaches to the same degree that a network made up of physical machines would be. Even if the virtual machine network is isolated from any physical network, virtual machines in the network can be subject to attacks from other virtual machines in the network. The requirements for securing virtual machines are often the same as those for physical machines.

Virtual machines are isolated from each other. One virtual machine cannot read or write another virtual machine's memory, access its data, use its applications, and so forth. However, within the network, any virtual machine or group of virtual machines can still be the target of unauthorized access from other virtual machines and might require further protection by external means. You can add this level of security by:

- Adding firewall protection to your virtual network by installing and configuring software firewalls on some or all of its virtual machines.

  For efficiency, you can set up private virtual machine Ethernet networks, or *virtual networks*. With virtual networks, you install a software firewall on a virtual machine at the head of the virtual network. This serves as a protective buffer between the physical network adapter and the remaining virtual machines in the virtual network.

  Installing a software firewall on virtual machines at the head of virtual networks is a good security practice. However, because software firewalls can slow performance, balance your security needs against performance before you decide to install software firewalls on virtual machines elsewhere in the virtual network. See "Networking Concepts Overview" on page 20.

- Keeping different virtual machine zones within a host on different network segments. If you isolate virtual machine zones on their own network segments, you minimize the risks of data leakage from one virtual machine zone to the next. Segmentation prevents various threats, including Address Resolution Protocol (ARP) spoofing, in which an attacker manipulates the ARP table to remap MAC and IP addresses, thereby gaining access to network traffic to and from a host. Attackers use ARP spoofing to generate denials of service, hijack the target system, and otherwise disrupt the virtual network.

Planning segmentation carefully lowers the chances of packet transmissions between virtual machine zones, thus preventing sniffing attacks, which require sending network traffic to the victim. Also, an attacker cannot use an insecure service in one virtual machine zone to access other virtual machine zones in the host. You can implement segmentation by using either of two approaches, each of which has different benefits.

- Use separate physical network adapters for virtual machine zones to ensure that the zones are isolated. Maintaining separate physical network adapters for virtual machine zones is probably the most secure method and is less prone to misconfiguration after the initial segment creation.
- Set up virtual local area networks (VLANs) to help safeguard your network. Because VLANs provide almost all of the security benefits inherent in implementing physically separate networks without the hardware overhead, they offer a viable solution that can save you the cost of deploying and maintaining additional devices, cabling, and so forth.

VLANs are an IEEE standard networking scheme with specific tagging methods that allow routing of packets to only those ports that are part of the VLAN. When properly configured, VLANs provide a dependable means for you to protect a set of virtual machines from accidental or malicious intrusions.

VLANs let you segment a physical network so that two machines in the network are unable to transmit packets back and forth unless they are part of the same VLAN. For example, accounting records and transactions are among a company's most sensitive internal information. In a company whose sales, shipping, and accounting employees all use virtual machines in the same physical network, you might protect the virtual machines for the accounting department by setting up VLANs as shown in Figure 10-4.

Figure 10-4. Sample VLAN Layout

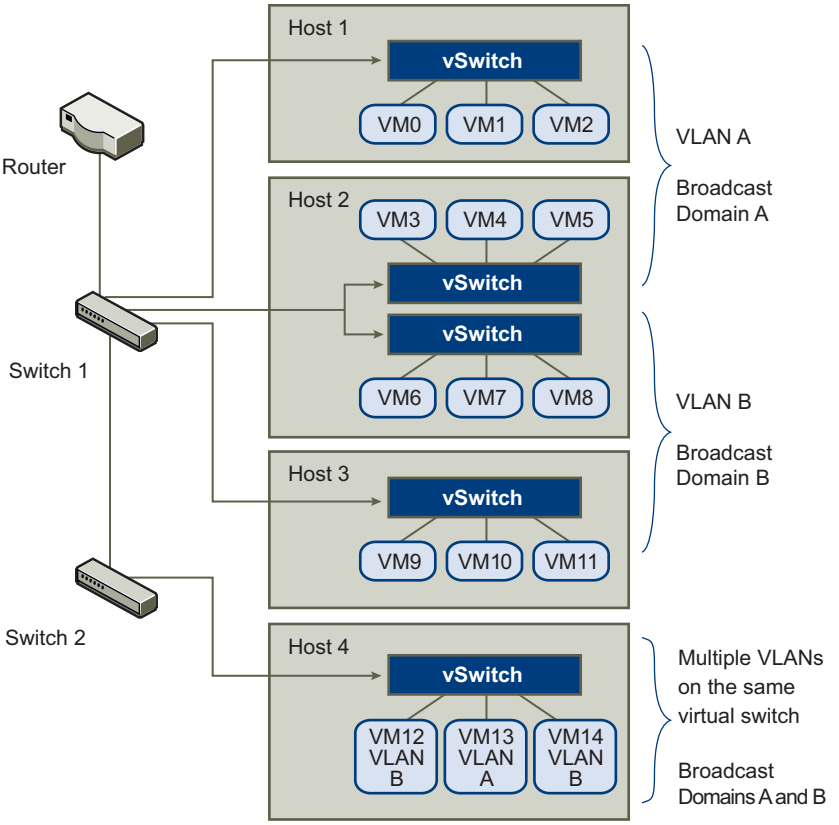

In this configuration, all employees in the accounting department use virtual machines in VLAN A and the employees in sales use virtual machines in VLAN B.

The router forwards packets containing accounting data to the switches. These packets are tagged for distribution to VLAN A only. Therefore, the data is confined to Broadcast Domain A, and cannot be routed to Broadcast Domain B unless the router is configured to do so.

This VLAN configuration prevents the sales force from intercepting packets destined for the accounting department. It also prevents the accounting department from receiving packets intended for the sales group. The virtual machines serviced by a single virtual switch can be in different VLANs.

Security Considerations for vSwitches and VLANs

ESX Server 3 features a complete IEEE 802.1q-compliant VLAN implementation. The way you set up VLANs to secure parts of a network depends on factors such as the guest operating system you install, the way your network equipment is configured, and so forth. While VMware cannot make specific recommendations on how to set up VLANs, consider the following factors when using a VLAN deployment as part of your security enforcement policy:

- **Treat VLANs as part of a broader security implementation** – VLANs are an effective means of controlling where and how widely data is transmitted within the network. If an attacker gains access to the network, the attack is likely to be limited to the VLAN that served as the entry point, lessening the risk to the network as a whole.

  VLANs provide protection only in that they control how data is routed and contained after it passes through the switches and enters the network. You can use VLANs to help secure Layer 2 of your network architecture—the data link layer. However, configuring VLANs doesn't protect the physical layer of your network model or any of the other layers. Even if you create VLANs, provide additional protection by securing your hardware (routers, hubs, and so forth) and encrypting data transmissions.

  VLANs are not a substitute for firewalls in your virtual machine configurations. Most network configurations that include VLANs also include software firewalls. If you include VLANs in your virtual network, be sure that any firewalls you install are VLAN-aware.

- **Be sure your VLANs are properly configured** – Equipment misconfiguration and network hardware, firmware, or software defects can make a VLAN susceptible to VLAN-hopping attacks. VLAN hopping occurs when an attacker with authorized access to one VLAN creates packets that trick physical switches into transmitting the packets to another VLAN that the attacker is not authorized to access. Vulnerability to this type of attack usually results from a switch being misconfigured for native VLAN operation, in which the switch can receive and transmit untagged packets.

  To help prevent VLAN hopping, keep your equipment up to date by installing hardware and firmware updates as they become available. Also, follow your vendor's best practice guidelines when you configure your equipment.

  VMware virtual switches do not support the concept of a native VLAN. All data passed on these switches is appropriately tagged. However, because other switches in the network might be configured for native VLAN operation, VLANs configured with virtual switches can still be vulnerable to VLAN hopping.

If you plan to use VLANs to enforce network security, VMware recommends that you disable the native VLAN feature for all switches unless you have a compelling need to operate some of your VLANs in native mode. If you need to use native VLAN, pay attention to your switch vendor's configuration guidelines for this feature.

- **Create a separate VLAN or virtual switch for communication between management tools and the service console** – Whether you use a management client or the command line, all configuration tasks for ESX Server 3 are performed through the service console, including configuring storage, controlling aspects of virtual machine behavior, and setting up virtual switches or virtual networks. Because the service console is the point of control for ESX Server 3, safeguarding it from misuse is crucial.

  VMware ESX Server 3 management clients use authentication and encryption to prevent unauthorized access to the service console, other services might not offer the same protection. If attackers gain access to the service console, they are free to reconfigure many attributes of the ESX Server 3 host. For example, they can change the entire virtual switch configuration, change authorization methods, and so forth.

  Network connectivity for the service console is established through virtual switches. To provide better protection for this critical ESX Server 3 component, VMware recommends that you isolate the service console by using one of the following methods:

  - Create a separate VLAN for management tool communication with the service console.
  - Configure network access for management tool connections with the service console through a single virtual switch and one or more uplink ports.

  Both methods prevent anyone without access to the service console VLAN or virtual switch from seeing traffic to and from the service console. They also prevent attackers from sending any packets to the service console. As an alternative, you can choose to configure the service console on a separate physical network segment instead. Physical segmentation provides a degree of additional security because it is less prone to later misconfiguration

  In addition to setting up a separate VLAN or virtual switch for management tool communication with the service console, you should set up separate a VLAN or virtual switch for VMotion and for network attached storage.

If your configuration includes an iSCSI SAN configured directly through the host rather than through a hardware adapter, you should create a separate virtual switch that provides shared network connectivity for the service console and for iSCSI. This second network connection for the service console is in addition to the primary service console network connection that you use for management tool communications. The second service console network connection supports iSCSI activities only, and you should not use it for any management activities or management tool communications.

Virtual Switch Protection and VLANs

VMware virtual switches provide safeguards against certain threats to VLAN security. Because of the way that virtual switches are designed, they protect VLANs against a variety of attacks, many of which involve VLAN hopping. Having this protection does not guarantee that your virtual machine configuration is invulnerable to other types of attacks. For example, virtual switches do not protect the physical network against these attacks, just the virtual network.

The following list gives you an idea of some attacks that virtual switches and VLANs can protect against.

- **MAC flooding** – Floods a switch with packets that contain MAC addresses tagged as having come from different sources. Many switches use a content-addressable memory (CAM) table to learn and store the source address for each packet. When the table is full, the switch can enter a fully open state in which every incoming packet is broadcast on all ports, letting the attacker see all of the switch's traffic. This state might result in packet leakage across VLANs.

  While VMware virtual switches store a MAC address table, they don't get the MAC addresses from observable traffic and are not vulnerable to this type of attack.

- **802.1q and ISL tagging attacks** – Force a switch to redirect frames from one VLAN to another by tricking the switch into acting as a trunk and broadcasting the traffic to other VLANs.

  VMware virtual switches don't perform the dynamic trunking required for this type of attack and, therefore, are not vulnerable.

- **Double-encapsulation attacks** – Occur when an attacker creates a double-encapsulated packet in which the VLAN identifier in the inner tag is different from the VLAN identifier in the outer tag. For backward compatibility, native VLANs strip the outer tag from transmitted packets unless configured to do otherwise. When a native VLAN switch strips the outer tag, only the inner tag is left, and that inner tag routes the packet to a different VLAN than the one identified in the now-missing outer tag.

  VMware virtual switches drop any double-encapsulated frames that a virtual machine attempts to send on a port configured for a specific VLAN. Therefore, they are not vulnerable to this type of attack.

- **Multicast brute-force attacks** – Involve sending large numbers of multicast frames to a known VLAN almost simultaneously in hopes of overloading the switch so that it mistakenly allows some of the frames to broadcast to other VLANs.

  VMware virtual switches do not allow frames to leave their correct broadcast domain (VLAN) and are not vulnerable to this type of attack.

- **Spanning-tree attacks** – Target Spanning-Tree Protocol (STP), which is used to control bridging between parts of the LAN. The attacker sends Bridge Protocol Data Unit (BPDU) packets that attempt to change the network topology, establishing themselves as the root bridge. As the root bridge, the attacker can sniff the contents of transmitted frames.

  VMware virtual switches don't support STP and are not vulnerable to this type of attack.

- **Random frame attacks** – Involve sending large numbers of packets in which the source and destination addresses stay the same, but in which fields are randomly changed in length, type, or content. The goal of this attack is to force packets to be mistakenly rerouted to a different VLAN.

  VMware virtual switches are not vulnerable to this type of attack.

Because new security threats develop over time, do not consider this an exhaustive list of attacks. Regularly check VMware security resources on the Web (http://www.vmware.com/support/security.html) to learn about security, recent security alerts, and VMware security tactics.

Securing Virtual Switch Ports

As with physical network adapters, a virtual network adapter can send frames that appear to be from a different machine or impersonate another machine so that it can receive network frames intended for that machine. Also, like physical network adapters, a virtual network adapter can be configured so that it receives frames targeted for other machines.

When you create a virtual switch for your network, you add port groups to impose a policy configuration for the virtual machines, storage systems, and so forth attached to the switch. You create virtual ports through the VI Client.

As part of adding a port or port group to a virtual switch, the VI Client configures a security profile for the port. You can use this security profile to ensure that ESX Server 3 prevents the guest operating systems for its virtual machines from impersonating other machines on the network. This security feature is implemented so that the guest operating system responsible for the impersonation does not detect that the impersonation was prevented.

The security profile determines how strongly you enforce protection against impersonation and interception attacks on virtual machines. To correctly use the settings in the security profile, you need to understand the basics of how virtual network adapters control transmissions and how attacks are staged at this level.

Each virtual network adapter has its own MAC address assigned when the adapter is created. This address is called the initial MAC address. Although the initial MAC address can be reconfigured from outside the guest operating system, it cannot be changed by the guest operating system. In addition, each adapter has an effective MAC address that filters out incoming network traffic with a destination MAC address different from the effective MAC address. The guest operating system is responsible for setting the effective MAC address and typically matches the effective MAC address to the initial MAC address.

When sending packets, an operating system typically places its own network adapter's effective MAC address in the source MAC address field of the Ethernet frame. It also places the MAC address for the receiving network adapter in the destination MAC address field. The receiving adapter accepts packets only when the destination MAC address in the packet matches its own effective MAC address.

Upon creation, a network adapter's effective MAC address and initial MAC address are the same. The virtual machine's operating system can alter the effective MAC address to another value at any time. If an operating system changes the effective MAC address, its network adapter then receives network traffic destined for the new MAC address. The operating system can send frames with an impersonated source MAC address at any time. Thus, an operating system can stage malicious attacks on the devices in a network by impersonating a network adapter that the receiving network authorizes.

You can use virtual switch security profiles on ESX Server 3 hosts to protect against this type of attack by setting the following options:

- **MAC address changes** – By default, this option is set to **Accept**, meaning that the ESX Server 3 host accepts requests to change the effective MAC address to other than the initial MAC address. The **MAC Address Changes** option setting affects traffic that a virtual machine receives.

  To protect against MAC impersonation, you can set this option to **Reject**. If you do, the ESX Server 3 host does not honor requests to change the effective MAC address to anything other than the initial MAC address. Instead, the port that the virtual adapter used to send the request is disabled. As a result, the virtual adapter does not receive any more frames until it changes the effective MAC address to match the initial MAC address. The guest operating system does not detect that the MAC address change was not honored.

  NOTE In some situations, you might have a legitimate need for more than one adapter to have the same MAC address on a network—for example, if you are using Microsoft Network Load Balancing in unicast mode. When Microsoft Network Load Balancing is used in the standard multicast mode, adapters do not share MAC addresses.

- **Forged transmissions** – By default, this option is set to **Accept**, meaning the ESX Server 3 host does not compare source and effective MAC addresses. The **Forged Trasmits** option setting affects traffic transmitted from a virtual machine.

  To protect against MAC impersonation, you can set this option to **Reject**. If you do, the ESX Server 3 host compares the source MAC address being transmitted by the operating system with the effective MAC address for its adapter to see if they match. If the addresses don't match, ESX Server 3 drops the packet.

  The guest operating system does not detect that its virtual network adapter cannot send packets by using the impersonated MAC address. The ESX Server 3 host intercepts any packets with impersonated addresses before they are delivered, and the guest operating system might assume that the packets are dropped.

- **Promiscuous mode operation** – By default, this option is set to **Reject**, meaning that the virtual network adapter cannot operate in promiscuous mode. Promiscuous mode eliminates any reception filtering that the virtual network adapter would perform, so that the guest operating system receives all traffic observed on the wire.

  Although promiscuous mode can be useful for tracking network activity, it is an insecure mode of operation because any adapter in promiscuous mode has access to the packets regardless of whether some of the packets should be received only by a particular network adapter. This means that an administrator or root user within a virtual machine can potentially view traffic destined for other guest or host operating systems.

  NOTE In some situations, you might have a legitimate need to configure a virtual switch to operate in promiscuous mode—for example, if you are running network intrusion detection software or a packet sniffer.

To change any of these default settings for a port, modify the security profile by editing virtual switch settings in the VI Client. For information on editing these settings, see "Virtual Switch Policies" on page 50.

Securing iSCSI Storage

The storage you configure for an ESX Server 3 host might include one or more storage area networks (SANs) that use iSCSI. iSCSI is a means of accessing SCSI devices and exchanging data records by using TCP/IP protocol over a network port rather than through a direct connection to a SCSI device. In iSCSI transactions, blocks of raw SCSI data are encapsulated in iSCSI records and transmitted to the requesting device or user.

iSCSI SANs let you make efficient use of existing Ethernet infrastructures to provide ESX Server 3 hosts access to storage resources that they can dynamically share. iSCSI SANs provide an economical storage solution for environments that rely on a common storage pool to serve numerous users. As with any networked system, your iSCSI SANs can be subject to security breaches. When you configure iSCSI on an ESX Server 3 host, you can take several measures to minimize security risks.

NOTE The requirements and procedures for securing an iSCSI SAN are similar for the hardware iSCSI adapters you can use with ESX Server 3 hosts and for iSCSI configured directly through the ESX Server 3 host. For information on configuring iSCSI adapters and storage, see "iSCSI Storage" on page 110.

Securing iSCSI Devices Through Authentication

One means of securing iSCSI devices from unwanted intrusion is to require that the ESX Server 3 host, or *initiator*, be authenticated by the iSCSI device, or *target*, whenever the host attempts to access data on the target LUN. The goal of authentication is to prove that the initiator has the right to access a target, a right granted when you configure authentication.

You have two choices when you set up authentication for iSCSI SANs on the ESX Server 3 host:

- **Challenge Handshake Authentication Protocol (CHAP)** – You can configure the iSCSI SAN to use CHAP authentication. In CHAP authentication, when the initiator contacts an iSCSI target, the target sends a predefined ID value and a random value, or *key*, to the initiator. The initiator then creates a one-way hash value that it sends to the target. The hash contains three elements: a predefined ID value, the random value that the target sends, and a private value, or *CHAP secret*, that the initiator and target share. When the target receives the hash from the initiator, it creates its own hash value by using the same elements and compares it to the initiator's hash. If the results match, the target authenticates the initiator.

  ESX Server 3 supports one-way CHAP authentication for iSCSI. It does not support bi-directional CHAP. In one-way CHAP authentication, the target authenticates the initiator, but the initiator does not authenticate the target. The initiator has only one set of credentials, and all of the iSCSI targets use them.

  ESX Server 3 supports CHAP authentication at the HBA level only. It does not support per-target CHAP authentication for each, which enables you to configure different credentials for each target to achieve greater target refinement.

- **Disabled** – You can configure the iSCSI SAN to use no authentication. Communications between the initiator and target are still authenticated in a rudimentary way because the iSCSI target devices are typically set up to communicate with specific initiators only.

  Choosing not to enforce more stringent authentication can make sense if your iSCSI storage is housed in one location and you create a dedicated network or VLAN to service all your iSCSI devices. The iSCSI configuration is secure because it is isolated from any unwanted access, much as a Fibre Channel SAN would be.

  As a basic rule, disable authentication only if you are willing to risk an attack to the iSCSI SAN or cope with problems that result from human error.

ESX Server 3 does not support Kerberos, Secure Remote Protocol (SRP), or public-key authentication methods for iSCSI. Additionally, it does not support IPsec authentication and encryption.

Use the VI Client to determine whether authentication is currently being performed and to configure the authentication method.

To check the authentication method

1 Log in to the VI Client and select the server from the inventory panel.

2 Click the **Configuration** tab and click **Storage Adapters**.

3 Select the iSCSI adapter to check and click **Properties** to open the **iSCSI Initiator Properties** dialog box.

4 Click **CHAP Authentication**.

If **CHAP Name** shows a name—often the iSCSI initiator name, the iSCSI SAN is using CHAP authentication.

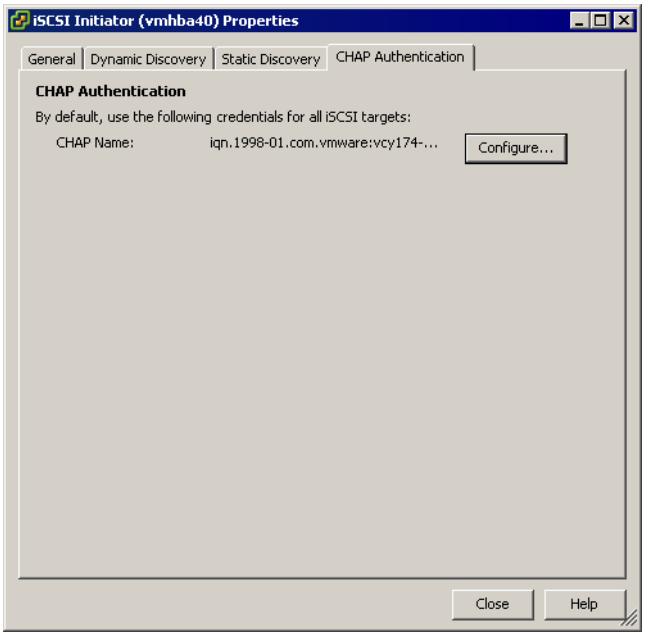

NOTE If **CHAP Name** shows **Not Specified**, the iSCSI SAN is not using CHAP authentication.

5 Click **Close**.

To configure iSCSI for CHAP authentication

1 Log in to the VI Client and select the server from the inventory panel.

2 Click the **Configuration** tab and click **Storage Adapters**.

3 Select the iSCSI adapter and click **Properties** to open the **iSCSI Initiator Properties** dialog box.

4 Click **CHAP Authentication** > **Configure** to open the **CHAP Authentication** dialog box.

5 Click **Use the following CHAP credentials**.

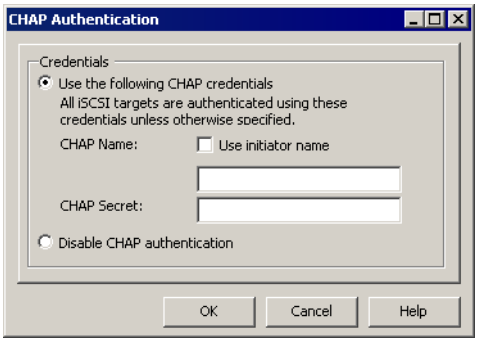

6 Perform one of the following actions:

- To set the CHAP name to the iSCSI adapter name, select **Use initiator name**.
- To set the CHAP name to anything other than the iSCSI adapter name, deselect **Use initiator name** and enter a name of up to 255 alphanumeric characters in the **CHAP Name** field.

7 Enter a CHAP secret to be used as part of authentication.

The secret you enter is a text string.

The VI Client doesn't impose a minimum or maximum length for the CHAP secret you enter. However, some iSCSI storage devices require that the secret exceed a minimum number of characters or have limitations on the character types you can use. To determine the requirements, check the manufacturer's documentation.

8 Click **OK**.

To disable iSCSI authentication

1 Log in to the VI Client and select the server from the inventory panel.

2 Click the **Configuration** tab and click **Storage Adapters**.

3 Select the iSCSI adapter and click **Properties** to open the **iSCSI Initiator Properties** dialog box.

4 Click **CHAP Authentication** > **Configure** to open the **CHAP Authentication** dialog box.

5 Select **Disable CHAP authentication**.

6 Click **OK**.

Protecting an iSCSI SAN

When you plan your iSCSI configuration, take measures to improve the overall security of the iSCSI SAN. Your iSCSI configuration is only as secure as your IP network, so by enforcing good security standards when you set up your network, you help safeguard your iSCSI storage.

The following are some specific suggestions:

- **Protect transmitted data** – A primary security risk in iSCSI SANs is that an attacker might sniff transmitted storage data.

  VMware recommends that you take additional measures to prevent attackers from easily seeing iSCSI data. Neither the hardware iSCSI adapter nor the ESX Server 3 host iSCSI initiator encrypts the data they transmit to and from the targets, making the data more vulnerable to sniffing attacks.

  Allowing your virtual machines to share virtual switches and VLANs with your iSCSI configuration potentially exposes iSCSI traffic to misuse by a virtual machine attacker. To help ensure that intruders can't listen to iSCSI transmissions, make sure that none of your virtual machines can see the iSCSI storage network.

  If you use a hardware iSCSI adapter, you can accomplish this by making sure that the iSCSI adapter and ESX Server 3 physical network adapter are not inadvertently connected outside the host by virtue of sharing a switch or some other means. If you configure iSCSI directly through the ESX Server 3 host, you can accomplish this by configuring iSCSI storage through a different virtual switch than the one used by your virtual machines, as shown in Figure 10-5.

Figure 10-5. iSCSI Storage on a Separate Virtual Switch

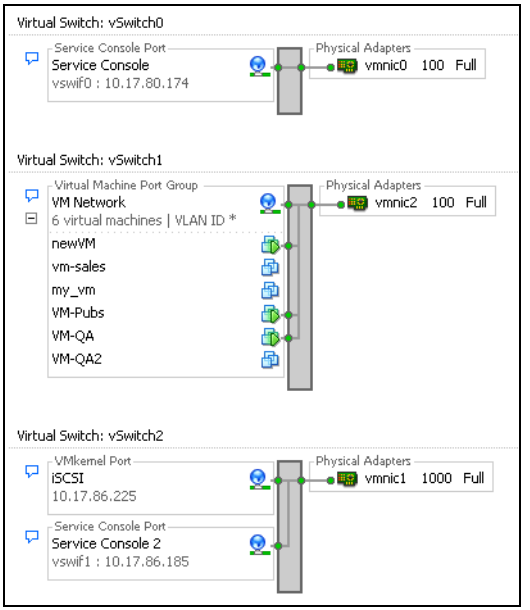

If you configure iSCSI directly through the host rather than through a hardware adapter, you must create two network connections for the service console within the virtual network setup. You configure the first service console network connection on its own virtual switch and use it exclusively for management tool connectivity (vSwitch0 in the figure). You configure the second service console network connection so that it shares the virtual switch you use for iSCSI connectivity (vSwitch2 in the figure). The second service console network connection supports iSCSI activities only. Do not use it for any management activities or management tool communications. To enforce a degree of separation between the iSCSI and service console on the shared virtual switch, configure them on different VLANs.

NOTE Do not configure the default gateway for the service console on the virtual switch you use for iSCSI connectivity. Configure it on the virtual switch you use for management tool connectivity instead.

In addition to protecting the iSCSI SAN by giving it a dedicated virtual switch, consider configuring your iSCSI SAN on its own VLAN. Placing your iSCSI configuration on a separate VLAN ensures that no devices other than the iSCSI adapter have visibility into transmissions within the iSCSI SAN.

- **Secure iSCSI ports** – When you run iSCSI devices, the ESX Server 3 host doesn't open any ports that listen for network connections. This measure reduces the chances that an intruder can break into the ESX Server 3 host through spare ports and gain control over the host. Therefore, running iSCSI doesn't present any additional security risks at the ESX Server 3 host end of the connection.

  Any iSCSI target device that you run must have one or more open TCP ports used to listen for iSCSI connections. If any security vulnerabilities exist in the iSCSI device software, your data can be at risk through no fault of ESX Server 3. To lower this risk, install all security patches that your storage equipment manufacturer provides and limit the devices connected to the iSCSI network.

Authentication and User Management

11

This chapter explains how ESX Server 3 handles user authentication and shows you how to set up user and group permissions. In addition, it discusses encryption for connections to the VI Client, SDK, and VI Web Access, as well as how to configure a delegate user name for transactions with NFS storage.

This chapter discusses the following topics:

- "Securing ESX Server 3 Through Authentication and Permissions" on page 211
- "Encryption and Security Certificates for ESX Server 3" on page 225
- "Virtual Machine Delegates for NFS Storage" on page 232

Securing ESX Server 3 Through Authentication and Permissions

ESX Server 3 uses the Pluggable Authentication Modules (PAM) structure for authentication when users access the ESX Server 3 host using the VI Client, VI Web Access, or the service console. The PAM configuration for VMware services is located in `/etc/pam.d/vmware-authd`, which stores paths to authentication modules.

The default installation of ESX Server 3 uses `/etc/passwd` authentication, just as Linux does, but you can configure ESX Server 3 to use another distributed authentication mechanism. If you plan to use a third-party authentication tool instead of the ESX Server 3 default implementation, see the vendor documentation for instructions. As part of setting up third-party authentication, you might need to update the `/etc/pam.d/vmware-authd` file with new module information.

Every time a VI Client or VirtualCenter user connects to an ESX Server 3 host, the `xinetd` process connects with the VMware Host Agent (`vmware-hostd`) process. The `vmware-hostd` process receives the user name and password from the client and forwards them to the PAM module to perform the authentication.

Figure 11-1 shows a basic example of how ESX Server 3 authenticates transactions from the VI Client.

Figure 11-1. Authentication for VI Client Communications with ESX Server 3

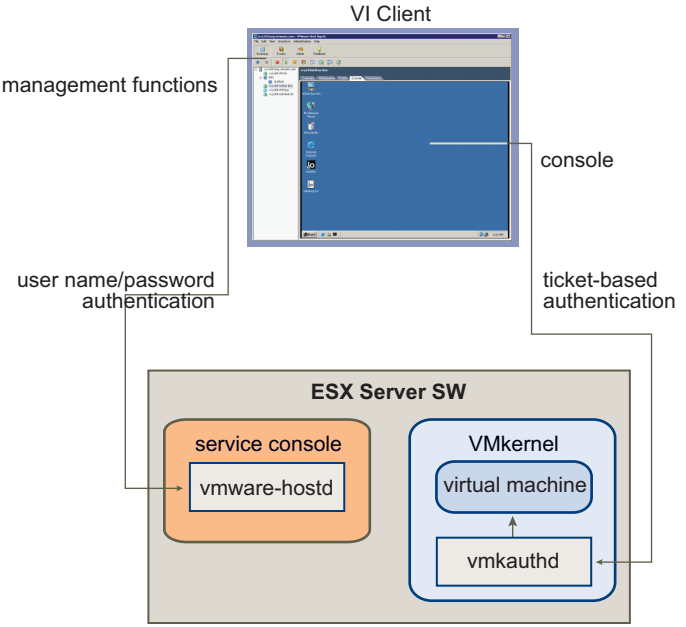

ESX Server authentication transactions with VI Web Access and third-party network management clients are similarly direct interactions with the `vmware-hostd` process.

To make sure that authentication works efficiently for your site, you might need to perform basic tasks such as setting up users, groups, permissions, and roles, configuring user attributes, adding your own certificates, determining whether you want to use SSL, and so forth.

About Users, Groups, Permissions, and Roles

Access to an ESX Server host and its resources is granted when a known user with appropriate permissions logs in to the host with a password that matches the password stored for that user. VirtualCenter uses a similar approach when determining whether to grant access to a user. VirtualCenter and ESX Server hosts determine the level of access for the user based on the permissions assigned to the user. For example, one user might have permissions that allows them to create virtual machines on the host and another user might have permissions that allow them to power on virtual machines but not create them.

The combination of user name, password, and permissions is the mechanism by which VirtualCenter and ESX Server 3 hosts authenticate a user for access and authorize the user to perform activities. To support this mechanism, the VirtualCenter and ESX Server hosts maintain lists of authorized users, their passwords, and the permissions assigned to each user. VirtualCenter and ESX Server hosts deny access under the following circumstances:

- A user not in the user list attempts to log in.
- A user enters the wrong password.
- A user is in the list but was not assigned permissions.
- A user who successfully logged in attempts operations that they do not have permission to perform.

As part of managing ESX Server hosts and VirtualCenter, you need to develop user and permission models, which are basic plans for how to handle particular types of users and how to design your permissions. In developing your user and permission models, be aware that:

- ESX Server 3 and VirtualCenter use sets of privileges, or *roles*, to control which operations individual users or groups can perform. ESX Server 3 and VirtualCenter provide you with a set of pre-established roles, but you can also create new ones.
- You can manage users more easily by assigning them to groups. If you create groups, you can apply a role to the group, and this role is inherited by all the users in the group.

Understanding Users

A user is an individual authorized to log in to either an ESX Server 3 host or to VirtualCenter. ESX Server 3 users fall into two categories: those who can access the ESX Server 3 host through VirtualCenter and those who can access the ESX Server 3 host by directly logging in to the host from the VI Client, VI Web Access, a third-party client, or a command shell. These two categories draw users from different sources.

- **Authorized VirtualCenter** – Are included in the Windows domain list that VirtualCenter references or are local Windows users on the VirtualCenter host.

  You cannot use VirtualCenter to manually create, remove, or otherwise change users. To manipulate the user list or change user passwords, you must do so through the tools you use to manage your Windows domain.

  Any changes you make to the Windows domain are reflected in VirtualCenter. However, because you cannot directly manage users in VirtualCenter, the user interface doesn't provide a user list for you to review. The only time you work with user and group lists is when you select users and groups during role assignment. You notice these changes only when you select users to configure permissions.

- **Direct-access users** – Authorized to work directly on an ESX Server 3 host are added to the internal user list by default when ESX Server 3 is installed or by a system administrator after installation.

  If you log in to the host as an administrator, you can perform a variety of management activities for these users, such as changing passwords, group memberships, permissions, and so forth. You can also add and remove users.

The user list that VirtualCenter maintains is completely separate from the user list that the ESX Server 3 host maintains. Even if the lists a host and VirtualCenter maintain appear to have common users (for instance, a user called *devuser*), treat these users as separate users who happen to have the same name. The attributes of devuser in VirtualCenter, including permissions, passwords, and so forth are separate from the attributes of devuser on the ESX Server 3 host. If you log in to VirtualCenter as devuser, you might have permission to view and delete files from a datastore, whereas if you log in to an ESX Server 3 host as devuser, you might not.

Because of the confusion that duplicate naming can cause, VMware recommends that you check the VirtualCenter user list before you create ESX Server 3 host users so that you can avoid duplicating names. To check for VirtualCenter users, review the Windows domain list.

Understanding Groups

You can more efficiently manage some user attributes by creating groups. A group is a set of users that you want to manage through a common set of rules and permissions. When you assign permissions to a group, all users in the group inherit them, and you do not have to work with the user profiles individually. Therefore, using groups can significantly reduce the time it takes to set up your permissions model and improve future scalability.

As an administrator, decide how to structure groups to achieve your security and usage goals. For example, three part-time sales team members work different days, and you want them to share a single virtual machine but not use the virtual machines belonging to sales managers. In this case, you might create a group called *SalesShare* that includes the three sales people: Mary, John, and Tom. You might then give the SalesShare group permission to interact with only one object, Virtual Machine A. Mary, John, and Tom inherit these permissions and are able to power up Virtual Machine A, start console sessions on Virtual Machine A, and so forth. They cannot perform these actions on the sales managers' virtual machines: Virtual Machines B, C, and D.

The group lists in VirtualCenter and an ESX Server 3 host are drawn from the same sources as their respective user lists. If you are working through VirtualCenter, the group list is called from the Windows domain. If you are logged in to an ESX Server 3 host directly, the group list is called from a table that the host maintains. All the recommendations for how you treat group lists are the same as those for user lists.

Understanding Permissions

For ESX Server 3 and VirtualCenter, permissions are defined as access roles that consist of a user and the user's assigned role for an object such as a virtual machine or ESX Server 3 host. Permissions grant users the right to perform specific activities and manage specific objects on an ESX Server 3 host or, if users are working from VirtualCenter, all VirtualCenter-managed objects. For example, to configure memory for an ESX Server 3 host, you must have a permission that grants host configuration privileges.

Most VirtualCenter and ESX Server 3 users have limited ability to manipulate the objects associated with the host. However, users in the Administrator role have full access rights and permissions on all virtual objects such as datastores, hosts, virtual machines, and resource pools. By default, the Administrator role is granted to the root user; if VirtualCenter manages the host, vpxuser is also an Administrator user. Administrator users have the following permissions:

- **root** – The root user can perform a complete range of control activities on the specific ESX Server 3 host that they are logged in to, including manipulating permissions, creating groups and users, working with events, and so forth. A root user logged in to one ESX Server 3 host cannot control the activities of any other host in the broader ESX Server 3 deployment.

  For security reasons, you might not want to use the root user in the Administrator role. In this case, you can change permissions after installation so that the root user no longer has administrative privileges or you can delete the root user's access permissions altogether through the VI Client as described in the "Managing Users, Groups, Permissions, and Roles" chapter of *Basic System Administration*. If you do so, you must first create another permission at the root level that has a different user assigned to the Administrator role.

  Assigning the Administrator role to a different user helps you maintain security through traceability. The VI Client logs all actions that the Administrator role user initiates as events, providing you with an audit trail. You can use this feature to improve accountability among the various users who act as administrators for the host. If all the administrators log in to the host as the root user, you cannot tell which administrator performed an action. If, instead, you create multiple permissions at the root level—each associated with a different user or user group—you can track the actions of each administrator or administrative group.

  After you create an alternative Administrator user, you can safely delete the root user's permissions or change its role to limit its privileges. If you delete or change the root user's permissions, you must use the new user you created as the host authentication point when you bring the host under VirtualCenter management. See "Understanding Roles" on page 218.

NOTE Configuration commands that you run through the command-line interface (`esxcfg` commands) do not perform an access check. Therefore, even if you limit the root user's privileges, this does not affect what that user can do using the command line interface commands.

- **vpxuser** – This user is VirtualCenter acting as an entity with Administrator rights on the ESX Server 3 host, allowing it to manage activities for that host. vpxuser is created at the time that an ESX Server 3 host is attached to VirtualCenter. It is not present on the ESX Server 3 host unless the host is being managed through VirtualCenter.

  When an ESX Server 3 host is managed through VirtualCenter, VirtualCenter has Administrator privileges on the host. For example, VirtualCenter can move virtual machines to and from hosts and perform configuration changes needed to support virtual machines.

  The VirtualCenter administrator, through vpxuser, can perform most of the same tasks on the host as the root user and also schedule tasks, work with templates, and so forth. However, you cannot perform certain activities as a VirtualCenter administrator. These activities, which include directly creating, deleting, or editing users and groups for ESX Server 3 hosts, can be performed only by a user with Administrator permissions directly on each ESX Server 3 host.

CAUTION Do not change vpxuser in any way and do not change its permissions. If you do so, you might experience problems in working with the ESX Server host through VirtualCenter.

If you are acting in the Administrator role on an ESX Server 3i host, you can grant permissions to individual users and groups on that host If you are acting in the Administrator role in VirtualCenter, you can grant permissions to any user or group included in the Windows domain list that VirtualCenter references.

VirtualCenter registers any selected Windows domain user or group through the process of assigning permissions. By default, all users who are members of the local Windows Administrators group on the VirtualCenter Server are granted the same access rights as any user assigned to the Administrator role. Users who are members of the Administrators group can log in as individuals and have full access.

For security reasons, consider removing the Windows Administrators group from the Administrator role. You can change permissions after installation so the Windows Administrators group does not have administrative privileges. Alternately you can use VI Client to delete the Windows Administrators group access permissions. If you delete Windows Administrators access permissions, you must first create another permission at the root level that has a different user assigned to the Administrator role.

The method you use to configure permissions directly on an ESX Server 3i host is identical to the method you use to configure permissions in VirtualCenter. Also, the list of privileges is the same for ESX Server 3i and VirtualCenter.

For information on configuring permissions and to read about the privileges you can assign, see *Basic System Administration*.

Understanding Roles

VirtualCenter and ESX Server grant access to objects only to users who are assigned permissions for the object. When you assign a user or group permissions for the object, you do so by pairing the user or group with a role. A role is a predefined set of privileges.

ESX Server hosts provide three default roles, and you cannot change the privileges associated with these roles. Each subsequent default role includes the privileges of the previous role. For example, the Administrator role inherits the privileges of the Read Only role. Roles you create yourself do not inherit privileges from any of the default roles. The default roles are:

- **No Access** – Users assigned this role for an object cannot view or change the object in any way. For example, a user who has a No Access role for a particular virtual machine cannot see the virtual machine in the VI Client inventory when they log in to the ESX Server host. With a No Access role for a particular object, a user can select the VI Client tabs associated with the no-access object, but the tab displays no content. For example, if the user does not have access to any virtual machines, they can select the **Virtual Machines** tab but will not see a virtual machine listing on the tab or any status information—the table is blank.

  The No Access role is the default assigned to any user or group you create on an ESX Server 3 host. You can elevate or lower a newly created user's or group's role on an object-by-object basis.

  The root user and vpxuser are the only users not assigned the No Access role by default. Instead, they are assigned the Administrator role.

  You can delete the root user's permissions altogether or change its role to No Access as long as you first create a replacement permission at the root level with the Administrator role and associate this role with a different user. If you delete or change the root user's permissions, you must use the new user you created as the host authentication point when you bring the host under VirtualCenter management.

- **Read Only** – Users assigned this role for an object are allowed to view the state of the object and details about the object.

  With this role, a user can view virtual machine, host, and resource pool attributes. The user cannot view the remote console for a host. All actions through the menus and toolbars are disallowed.

- **Administrator** – Users assigned this role for an object are allowed to view and perform all actions on the object. This role also includes all permissions inherent in the Read Only role.

You can create custom roles by using the role-editing facilities in the VI Client to create privilege sets that match your user needs. If you use the VI Client connected to VirtualCenter to manage your ESX Server 3 hosts, you have additional roles to choose from in VirtualCenter. Also, the roles you create directly on an ESX Server 3 host are not accessible within VirtualCenter. You can work with these roles, only if you log in to the host directly from the VI Client.

If you manage ESX Server 3 hosts through VirtualCenter, maintaining custom roles in the host and VirtualCenter can result in confusion and misuse. In this type of configuration, VMware recommends that you maintain custom roles only in VirtualCenter. For information on creating, altering, and deleting roles as well as a discussion of additional roles available in VirtualCenter, see *Basic System Administration*.

Working with Users and Groups on ESX Server 3 Hosts

If you are directly connected to an ESX Server 3 host through the VI Client, you can create, edit, and delete users and groups. These users and groups are visible in the VI Client whenever you log in to the ESX Server 3 host, but are not available if you log in to VirtualCenter.

Viewing and Exporting Users and Group Information

You work with users and groups through the **Users & Groups** tab in the VI Client. This tab displays a **Users** table or **Groups** table depending on whether you click **Users** or **Groups**.

You can also create roles and set permissions through a direct connection to the ESX Server 3 host. Because these tasks are more widely performed in VirtualCenter, see *Basic System Administration* for information on working with permissions and roles.

Figure 11-2 shows the **Users** table. The **Groups** table is similar.

Figure 11-2. Users Table

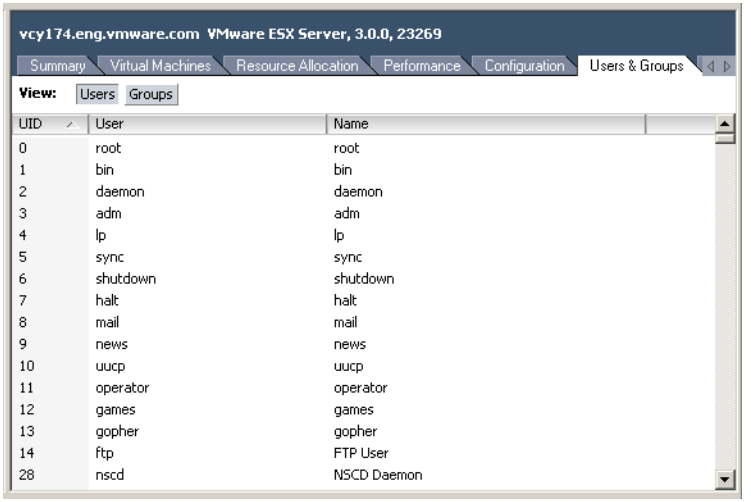

You can sort the lists according to column, show and hide columns, and export the list in formats you can use when preparing reports or publishing user or group lists on the Web.

To view and sort ESX Server 3 users or groups

1 Log in to the VI Client through the ESX Server 3 host.

2 Select the server from the inventory panel.

3 Click the **Users & Groups** tab and click **Users** or **Groups**.

4 Perform any of these actions as appropriate:

- To sort the table by any of the columns, click the column heading.
- To show or hide columns, right-click any of the column headings and select or deselect the name of the column to hide.

To export data in the ESX Server 3 Users or Groups table

1 Log in to the VI Client through the ESX Server 3 host.

2 Select the server from the inventory panel.

3 Click the **Users & Groups** tab and click **Users** or **Groups**.

4 Determine how to sort the table, and hide or show columns according to the information you want to see in the exported file.

5 Right-click anywhere in the user table and click **Export** to open the **Save As** dialog box.

6 Select a path, enter a filename, and select the file type.

7 Click **OK**.

Working with the Users Table

You can add users to the **Users** table for an ESX Server 3 host, remove users, and change various user attributes such as password and group memberships. When you perform these activities, you are altering the internal user list the ESX Server 3 host maintains.

To add a user to the ESX Server 3 Users table

1 Log in to the VI Client through the ESX Server 3 host.

2 Select the server from the inventory panel.

3 Click the **Users & Groups** tab and click **Users**.

4 Right-click anywhere in the **Users** table and click **Add** to open the **Add New User** dialog box.

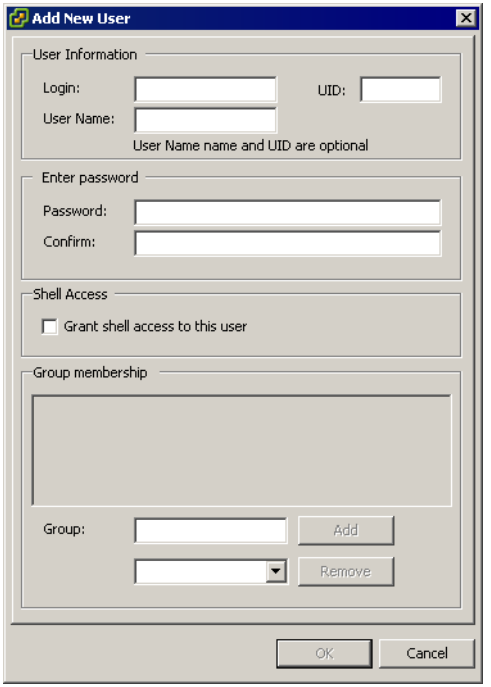

5 Enter a login, a user name, a numeric user ID (UID), and a password.

Specifying the user name and UID are optional. If you don't specify the UID, the VI Client assigns the next available UID.

The password should meet the length and complexity requirements outlined in "Password Restrictions" on page 241. However, the ESX Server 3 host checks for password compliance only if you have switched to the `pam_passwdqc.so` plug-in for authentication. The password settings in the default authentication plug-in, `pam_cracklib.so`, are not enforced.

6 If you want the user to be able to access the ESX Server 3 host through a command shell, select **Grant shell access to this user**.

In general, do not grant shell access to ESX Server 3 host users unless you determine that they have a justifiable need to access the host through a shell rather than through the VI Client. Users that access the host only through the VI Client do not need shell access.

7 For each existing group you want the user to be part of, enter the group name and click **Add**.

If you type a nonexistent group name, the VI Client warns you and does not add the group to the **Group membership** list.

8 Click **OK**.

The login and user name you entered now appear in the **Users** table.

To modify the settings for a user

1 Log in to the VI Client through the ESX Server 3 host.

2 Select the server from the inventory panel.

3 Click the **Users & Groups** tab and click **Users**.

4 Right-click the user to modify and click **Edit** to open the **Edit User** dialog box.

5 To change the user ID, enter a numeric user UID in the **UID** field.

The VI Client assigns the UID when you first create the user. In most cases, this assignment does not need to be changed.

6 Enter a new user name.

7 To change the user's password, select **Change Password** and enter the new password.

The password should meet the length and complexity requirements outlined in "Password Restrictions" on page 241. However, the ESX Server 3 host checks for password compliance only if you switched to the `pam_passwdqc.so` plug-in for authentication. The password settings in the default authentication plugin, `pam_cracklib.so`, are not enforced.

8 To change the user's ability to access the ESX Server 3 host through a command shell, select or deselect **Grant shell access to this user**.

9 To add the user to another group, enter the group name and click **Add**.

If you type a nonexistent group name, the VI Client warns you and does not add the group to the **Group membership** list.

10 To remove the user from a group, select the group name from the list and click **Remove**.

11 Click **OK**.

To remove a user from the ESX Server 3 Users table

1 Log on to the VI Client through the ESX Server 3 host.

2 Select the server from the inventory panel.

3 Click the **Users & Groups** tab and click **Users**.

4 Right-click the user to remove and click **Remove**.

CAUTION Do *not* remove the root user.

Working with the Groups Table

You can add groups to the **Groups** table for an ESX Server 3 host, remove groups, and add or remove group members. When you perform these activities, you are altering the internal group list that the ESX Server 3 host maintains.

To add a group to the ESX Server 3 Groups table

1 Log in to the VI Client through the ESX Server 3 host.

2 Select the server from the inventory panel.

3 Click the **Users & Groups** tab and click **Groups**.

4 Right-click anywhere in the **Groups** table and click **Add** to open the **Create New Group** dialog box.

5 Enter a group name and numeric group ID (GID) in the **Group ID** field.

Specifying the GID is optional. If you don't specify a GID, the VI Client assigns the next available group ID.

6 For each user that you want as a group member, enter the user name and click **Add**.

If you type a nonexistent user name, the VI Client warns you and does not add the user to the **Users in this group** list.

7 Click **OK**.

The group ID and group name you entered now appear in the **Groups** table.

To add or remove users from a group

1 Log in to the VI Client through the ESX Server 3 host.

2 Select the server from the inventory panel.

3 Click the **Users & Groups** tab and click **Groups**.

4 Right-click the group to modify and click **Edit** to open the **Edit Group** dialog box.

5 To add a user to the group, enter the user name and click **Add**.

If you type a nonexistent user name, the VI Client warns you and does not add the user to the **Users in this group** list.

6 To remove a user from the group, select the user name from the list and click **Remove**.

7 Click **OK**.

To remove a group from the ESX Server 3 Groups table

1 Log in to the VI Client through the ESX Server 3 host.

2 Select the server from the inventory panel.

3 Click the **Users & Groups** tab and click **Groups**.

4 Right-click the group you want to remove and click **Remove**.

CAUTION Do *not* remove the root user.

Encryption and Security Certificates for ESX Server 3

ESX Server supports SSL v3 and TLS v1, generally referred to here as SSL. SSL helps secure communications. If SSL is enabled data is private, so it can not be read, and it is protected, so it cannot be modified in transit without detection. All network traffic is encrypted as long as the following conditions are true:

- You did not change the Web proxy service to allow unencrypted traffic for the port.
- Your service console firewall is configured for medium or high security. For information on configuring the service console firewall, see "Service Console Firewall Configuration" on page 237.

(SEE UPDATE) SSL is not enabled by default, so network traffic is not encrypted unless you take action. SSL does protect the initial connection between VI Clients and VirtualCenter, but subsequent communications are not encrypted. To fully enable the security provided by certificates in ESX Server 3, you must enable certificate checking and install new certificates.

To enable certificate checking

1 Log in to a VirtualCenter server using the VI Client.

2 Click **Administration** > **Virtual Center Management Server Configuration**.

The **Virtual Center Management Server Configuration** dialog appears.

3 Click **SSL Settings** in the left pane and enable the **Check host certificates** checkbox.

4 Click **OK**.

(SEE UPDATE) To receive the full benefit of certificate checking, install new certificates. The initial certificates are created by ESX Server and stored on the host. The certificates used to secure your VirtualCenter and VI Web Access sessions are not signed by a trusted certificate authority and, therefore, do not provide the authentication security you might need in a production environment. For example, self-signed certificates are vulnerable to man-in-the-middle attacks. If you intend to use encrypted remote connections externally, consider purchasing a certificate from a trusted certificate authority or use your own security certificate for your SSL connections. If the self-signed certificate is used, clients receive a warning about the certificate.

To address this issue, install a certificate that is signed by a recognized certificate authority.

The default location for your certificate is `/etc/vmware/ssl/` on the ESX Server 3 host. The certificate consists of two files: the certificate itself (`rui.crt`) and the private-key file (`rui.key`).

Adding Certificates and Modifying ESX Server 3 Web Proxy Settings

When you add certificates for ESX Server 3, think about encryption and user security, be aware of the following:

- (SEE UPDATE) Avoid setting up certificates by using pass phrases. ESX Server 3 does not handle pass phrases, also known as encrypted keys. If you set up a pass phrase, ESX Server 3 processes cannot start correctly.
- You can configure the Web proxy so that it searches for certificates in a location other than the default location. This capability proves useful for companies that prefer to centralize their certificates on a single machine so that multiple hosts can use the certificates.

CAUTION Certificates stored in a location other the ESX Server 3 host are unusable if the host loses network connectivity. If certificate checking is enabled, you cannot establish secure connections with guests.

- To support encryption for user names, passwords, and packets, SSL is enabled by default for VI Web Access and VMware Infrastructure SDK connections. To configure these connections so that they don't encrypt transmissions, disable SSL for your VI Web Access connection or VMware Infrastructure SDK connection by switching the connection from HTTPS to HTTP, as described in "To change security settings for a Web proxy service" on page 228. Consider disabling SSL only if you created a fully trusted environment for these clients, where firewalls are in place and transmissions to and from the host are fully isolated. Disabling SSL can improve performance, because you avoid the overhead required to perform encryption.
- To protect against misuse of ESX Server 3 services such as the internal Web server that hosts VI Web Access, most internal ESX Server 3 services are accessible only through port 443, the port used for HTTPS transmission. Port 443 acts as a reverse proxy for ESX Server 3. You can see a list of services on ESX Server 3 through an HTTP welcome page, but you can't directly access these services without proper authorization. You can change this configuration so that individual services are directly accessible through HTTP connections. VMware recommends that you not make this change unless you are using ESX Server 3 in a fully trusted environment.
- When you upgrade VirtualCenter and VI Web Access, the certificate remains in place. If you remove VirtualCenter and VI Web Access, the certificate directory is not removed from the service console. (SEE UPDATE)

To configure the Web proxy to search for certificates in nondefault locations

1 Log in to the service console as the root user.

2 Change directories to `/etc/vmware/hostd/`.

3 Use a text editor to open the `proxy.xml` file and find the following XML segment:

```
<ssl>
   <!-- The server private key file -->
   <privateKey>/etc/vmware/ssl/rui.key</privateKey>
   <!-- The server side certificate file -->
   <certificate>/etc/vmware/ssl/rui.crt</certificate>
</ssl>
```

4 Replace `/etc/vmware/ssl/rui.key` with the absolute path to the private key file that you received from your trusted certificate authority.

This path can be on the ESX Server 3 host or on a centralized machine on which you store certificates and keys for your company.

NOTE Leave the `<privateKey>` and `</privateKey>` XML tags in place.

5 Replace `/etc/vmware/ssl/rui.crt` with the absolute path to the certificate file that you received from your trusted certificate authority.

CAUTION Do not delete the original `rui.key` and `rui.crt` files. The ESX Server host uses these files.

6 Save your changes and close the file.

7 Enter the following command to restart the `vmware-hostd` process:

```
service mgmt-vmware restart
```

To change security settings for a Web proxy service

1 Log in to the service console as the root user.

2 Change directories to `/etc/vmware/hostd/`.

3 Use a text editor to open the `proxy.xml` file. Contents of the file typically appears as follows:

```
<ConfigRoot>
   <EndpointList>
      <_length>6</_length>
      <_type>vim.ProxyService.EndpointSpec[]</_type>
      <e id="0">
         <_type>vim.ProxyService.NamedPipeServiceSpec</_type>
         <accessMode>httpsWithRedirect</accessMode>
         <pipeName>/var/run/vmware/proxy-webserver</pipeName>
         <serverNamespace>/</serverNamespace>
      </e>
      <e id="1">
         <_type>vim.ProxyService.NamedPipeServiceSpec</_type>
         <accessMode>httpsWithRedirect</accessMode>
         <pipeName>/var/run/vmware/proxy-sdk</pipeName>
         <serverNamespace>/sdk</serverNamespace>
      </e>
      <e id="2">
         <_type>vim.ProxyService.LocalServiceSpec</_type>
         <accessMode>httpsWithRedirect</accessMode>
         <port>8080</port>
         <serverNamespace>/ui</serverNamespace>
      </e>
      <e id="3">
         <_type>vim.ProxyService.NamedPipeServiceSpec</_type>
         <accessMode>httpsOnly</accessMode>
         <pipeName>/var/run/vmware/proxy-vpxa</pipeName>
         <serverNamespace>/vpxa</serverNamespace>
      </e>
      <e id="4">
         <_type>vim.ProxyService.NamedPipeServiceSpec</_type>
         <accessMode>httpsWithRedirect</accessMode>
         <pipeName>/var/run/vmware/proxy-mob</pipeName>
         <serverNamespace>/mob</serverNamespace>
      </e>
      <e id="5">
         <_type>vim.ProxyService.LocalServiceSpec</_type>
         <!-- Use this mode for "secure" deployment -->
         <!-- <accessMode>httpsWithRedirect</accessMode> -->
         <!-- Use this mode for "insecure" deployment -->
         <accessMode>httpAndHttps</accessMode>
         <port>8889</port>
         <serverNamespace>/wsman</serverNamespace>
      </e>
   </EndpointList>
</ConfigRoot>
```

4 Change the security settings as required. For example, you might want to modify entries for services that use HTTPS to add the option of HTTP access.

- *e id* is an ID number for the server ID XML tag. ID numbers must be unique within the HTTP area.
- *_type* is the name of the service you are moving, for example `/sdk` or `/mob`.
- *accessmode* is the forms of communication the service permits. Acceptable values include:
  - httpOnly – The service is accessible only over plain-text HTTP connections.
  - httpsOnly – The service is accessible only over HTTPS connections.
  - httpsWithRedirect – The service is accessible only over HTTPS connections. Requests over HTTP will be redirected to the appropriate HTTPS URL.
  - httpAndHttps – The service is accessible both over HTTP and HTTPS connections.
- *port* is the port number assigned to the service. You can assign a different port number to the service.
- *namespace* is the namespace for the server that provides this service.

5 Save your changes and close the file.

6 Enter the following command to restart the `vmware-hostd` process:

```
service mgmt-vmware restart
```

Example: Setting Up VI Web Access to Communicate Through an Insecure Port

VI Web Access normally communicates with an ESX Server 3 host through a secure port (HTTPS, 443). If you are in a fully trusted environment, you might decide that you can almost permit an insecure port (for example, HTTP, 80). To do so, change the accessMode attribute for the webserver in `proxy.xml` file, as described in "To change security settings for a Web proxy service" on page 228. The result is as follows. The accessMode is changed from `httpsWithRedirect` to `httpAndHttps`.

```
<_type>vim.ProxyService.LocalServiceSpec</_type>
<accessMode>httpAndHttps</accessMode>
<port>8080</port>
<serverNamespace>/ui</serverNamespace>
```

Regenerating Certificates

The ESX Server 3 host generates certificates when you first start the host after installation. Thereafter, whenever you restart the `vmware-hostd` process, the `mgmt-vmware` script searches for existing certificate files (`rui.crt` and `rui.key`) and, if it cannot find them, generates new certificate files.

Under certain circumstances, you might need to force the ESX Server 3 host to generate new certificates. You typically need to generate new certificates only if:

- You change the host name.
- You accidentally delete the certificates.

To generate new certificates for the ESX Server 3 host

1 Change directories to `/etc/vmware/ssl`.

2 Create backups of any existing certificates by executing the following commands:

```
mv rui.crt orig.rui.crt
mv rui.key orig.rui.key
```

NOTE If you are regenerating certificates because you accidentally deleted your certificates, you do not need to complete the backup step.

3 Enter the following command to restart the `vmware-hostd` process:

```
service mgmt-vmware restart
```

4 Confirm that the ESX Server 3 host generated new certificates by executing the following command comparing the time stamps of the new certificate files with orig.rui.crt and orig.rui.key:

```
ls -la
```

Replacing Self-Signed Certificates with CA-Signed Certificates

The ESX Server 3 host uses automatically generated self-signed certificates that are created as part of the installation process. These certificates make it possible to begin using the server, but self-signed certificates are vulnerable to a man-in-the-middle attack. To provide increased security, purchase a certificate from a trusted certificate authority and use that certificate as a replacement for the automatically generated certificate.

To replace the existing certificate with a CA-signed certificate

1 Change directories to `/etc/vmware/ssl`.

2 Backup any existing certificates by renaming them using the following commands:

```
mv rui.crt orig.rui.crt
mv rui.key orig.rui.key
```

3 Copy the new certificate and key to the current location. Rename the certificate and key to `rui.crt` and `rui.key`.

Virtual Machine Delegates for NFS Storage

To perform most activities on virtual machines, an ESX Server 3 needs access to virtual machine files. For instance, to power on and off virtual machines, ESX Server 3 must be able to create, manipulate, and delete files on the volume that is storing the virtual disk files.

To create, configure, or administer virtual machines on an NFS datastore, you can do so using the delegate user. The delegate user's identity is used by the ESX Server 3i for all I/O requests issued to the underlying file system. The delegate user is experimental and is not officially supported.

By default, the delegate user for the ESX Server 3 host is the user `root`. However, having the user `root` as the delegate user might not work for all NFS datastores. NFS administrators might export volumes with root squashing enabled. The `root squash` feature maps the user root to a user with no significant privileges on the NFS server, limiting the root user's abilities. This feature is commonly used to prevent unauthorized access to files on an NFS volume. If the NFS volume was exported with `root squash` enabled, the NFS server might refuse access to the ESX Server 3 host. To ensure that you can create and manage virtual machines from your host, the NFS administrator must turn off the `root squash` feature or add the ESX Server 3 host's physical network adapter to the list of trusted servers.

If the NFS administrator is unwilling to take either of these actions, you can change the delegate user to a different identity through experimental ESX Server 3 functionality. This identity must match the owner of the directory on the NFS server, or the ESX Server 3 host cannot perform file-level operations. To set up a different identity for the delegate user, acquire the following information:

- User name of the directory owner
- User ID (UID) of the directory owner
- Group ID (GID) of the directory owner

Use this information to change the delegate user setting for the ESX Server 3 host so that it matches the owner of the directory, enabling NFS datastore to recognize the ESX Server 3 host correctly. The delegate user is configured globally, and the same identity is used to access every volume.

Setting up the delegate user on an ESX Server 3 host requires that you complete these activities:

- From the **Users & Groups** tab for a VI Client running directly on the ESX Server 3 host, either:
  - Edit the user named `vimuser` to add the correct UID and GID. `vimuser` is an ESX Server 3 host user provided to you as a convenience for setting up delegate users. By default, `vimuser` has a UID of 12 and a GID of 20.
  - Add a completely new user to the ESX Server 3 host with the delegate user name, UID, and GID.

  You must perform one of these steps regardless of whether you manage the host through a direct connection or through the VirtualCenter Server. Also, make sure that the delegate user (`vimuser` or a delegate user you create) is identical across all ESX Server 3 hosts that use the NFS datastore. For information on adding users, see "Working with the Users Table" on page 221.
- Configure a virtual machine delegate as part of the security profile for the host, as described in the procedure that follows. You configure the security profile through VirtualCenter or through a VI Client running directly on the ESX Server 3 host.

WARNING Changing the delegate user for an ESX Server 3 host is experimental and, currently, VMware does not support this implementation. Use of this function can result in unexpected behavior.

To change the virtual machine delegate

1. Log in to the VI Client through the ESX Server 3 host.
2. Select the server from the inventory panel.

   The hardware configuration page for this server appears with the **Summary** tab displayed.
3. Click **Enter Maintenance Mode**.
4. Click the **Configuration** tab and click **Security Profile**.
5. Click **Virtual Machine Delegate > Edit** to open the **Virtual Machine Delegate** dialog box.

6 Enter the user name for the delegate user.

7 Click **OK**.

8 Reboot the ESX Server 3 host.

After you reboot the host, the delegate user setting is visible in VirtualCenter and in the VI Client running directly on the ESX Server 3 host.

12 Service Console Security

This chapter makes basic security recommendations for using the service console and explains some of the service console's built-in security features. The service console is a management interface to ESX Server 3 and, as such, its security is critical. To protect the service console against unauthorized intrusion and misuse, VMware imposes constraints on several service console parameters, settings, and activities.

These constraints are designed to raise the security level for ESX Server 3. You can loosen the constraints to meet your particular configuration needs, but if you do so, make sure you are working in a trusted environment and have taken enough other security measures to protect the network as a whole and the devices connected to the ESX Server 3 host.

This chapter discusses the following topics:

- "General Security Recommendations" on page 236
- "Service Console Firewall Configuration" on page 237
- "Password Restrictions" on page 241
- "Cipher Strength" on page 249
- "setuid and setgid Applications" on page 250
- "SSH Security" on page 253
- "Security Patches and Security Vulnerability Scanning Software" on page 254

General Security Recommendations

Consider the following recommendations when evaluating service console security and administering the service console:

- *Limit user access.*

  To improve security, restrict user access to the service console and enforce access security policies like setting up password restrictions—for example, character length, password aging limits, and using a `grub` password for booting the host.

  The service console has privileged access to certain parts of ESX Server 3. Therefore, only trusted users should be provided login access. By default, root access is limited by not allowing secure shell (SSH) login as the root user, and you should strongly consider keeping this default. ESX Server 3 system administrators should be required to log in as regular users and then use the `sudo` command to perform specific tasks that require root privileges.

  Also, try to run as few processes on the service console as possible. Ideally, strive to run only the essential processes, services, and agents such as virus checkers, virtual machine backups, and so forth.

- *Use VI Client to administer your ESX Server 3 hosts.*

  Whenever possible, use VI Client, VI Web Access, or a third-party network management tool to administer your ESX Server 3 hosts instead of working though the command-line interface as the root user. Using VI Client lets you limit the accounts with access to the service console, safely delegate responsibilities, and set up roles that prevent administrators and users from using capabilities they don't need.

- *Use only VMware sources to upgrade ESX Server 3 components that you run on the service console.*

  The service console runs a variety of third-party packages, such as the Tomcat Web service, to support management interfaces or tasks that you need to perform. VMware does not support upgrading these packages from anything other than a VMware source. If you use a download or patch from another source, you might compromise service console security or functions. Regularly check third-party vendor sites and the VMware knowledge base for security alerts.

Logging On to the Service Console

Although you perform most ESX Server 3 configuration activities through the VI Client, you use the service console command-line interface when you configure certain security features. Using the command-line interface requires that you log in to the host. If you have direct access to the ESX Server 3 host, you can log in to the physical console on that machine. To do so, press Alt+F2 to open the login page. For remote connections, use SSH or another remote console connection to start a session on the host.

Whether you access the service console locally or through a remote connection such as SSH, you must log in using a user name and password recognized by the ESX Server 3 host. For information on user names and passwords for ESX Server 3 hosts, see "Working with Users and Groups on ESX Server 3 Hosts" on page 219.

If you are logging onto the host to perform activities that require root privileges, you should log in to the service console as a recognized user and acquire root privileges through the `su` command or, preferably, the `sudo` command. The `sudo` command enhances security because it grants root privileges only for select activities in contrast to the `su` command, which grants root privileges for all activities. Using `sudo` also provides superior accountability because all `sudo` activities are logged, whereas if you use `su`, ESX Server 3 only logs the fact that the user switched to root by way of `su`.

In addition to ESX-specific commands, you can use the service console command-line interface to execute many Linux and Unix commands. For detailed usage notes on service console commands, use the `man <command_name>` command to check for man pages.

Service Console Firewall Configuration

ESX Server 3 includes a firewall between the service console and the network. To ensure the integrity of the service console, VMware has reduced the number of firewall ports that are open by default. At installation time, the service console firewall is configured to block all incoming and outgoing traffic except for traffic on ports 902, 80, 443, and 22, which are used for basic communication with ESX Server 3. This setting enforces a high level of security for the ESX Server 3 host.

NOTE The firewall also allows Internet Control Message Protocol (ICMP) pings and communication with DHCP and DNS (UDP only) clients.

In trusted environments, you might decide that a lower security level is acceptable. If so, you can set the firewall for either medium or low security:

- **Medium security** – All incoming traffic is blocked except on the default ports (902, 433, 80, and 22) and any ports you specifically open. Outgoing traffic is not blocked.
- **Low security** – There are no blocks on either incoming or outgoing traffic. This setting is equivalent to removing the firewall.

Because the ports open by default are strictly limited, you might need to open additional ports after installation. For a list of commonly used ports that you might need to open, see "TCP and UDP Ports for Management Access" on page 183.

As you add the supported services and management agents required to operate ESX Server 3 effectively, you open other ports in the service console firewall. You add services and management agents through VirtualCenter as described in "Opening Firewall Ports for Supported Services and Management Agents" on page 188.

In addition to the ports you open for these services and agents, you might need to open other ports when you configure certain devices, services, or agents such as storage devices, backup agents, and management agents. For example, if you are using Veritas NetBackup™ 4.5 as a backup agent, open ports 13720, 13724, 13782, and 13783, which NetBackup uses for client-media transactions, database backups, user backups or restores, and so forth. To determine which ports to open, see vendor specifications for the device, service, or agent.

Changing the Service Console Security Level

Altering the security level for the service console is a two-part process: determining the service console firewall security level and resetting the service console firewall setting. To prevent unnecessary steps, always check the firewall setting before you change it.

Each time you lower your security setting or open additional ports, you increase the risk of intrusion in your network. Balance your access needs against how tightly you want to control the security of the network.

To determine the service console firewall security level

1 Log in to the service console and acquire root privileges.

2 Execute the following two commands to determine whether incoming and outgoing traffic is blocked or allowed:

```
esxcfg-firewall -q incoming
esxcfg-firewall -q outgoing
```

3 Interpret the results as follows:

Table 12-1. Service Console Security Levels

Command Line Response	Security Level
`Incoming ports blocked by default.` `Outgoing ports blocked by default.`	High
`Incoming ports blocked by default.` `Outgoing ports not blocked by default.`	Medium
`Incoming ports not blocked by default.` `Outgoing ports not blocked by default.`	Low

To set the service console firewall security level

1 Log in to the service console and acquire root privileges.

2 Execute one of the following commands as applicable:

- To set the service console firewall to medium security:

  ```
  esxcfg-firewall --allowOutgoing --blockIncoming
  ```

- To set the virtual firewall to low security:

  ```
  esxcfg-firewall --allowIncoming --allowOutgoing
  ```

CAUTION Using the preceding command disables all firewall protection.

- To return the service console firewall to high security:

  ```
  esxcfg-firewall --blockIncoming --blockOutgoing
  ```

3 Execute the following command to restart the `vmware-hostd` process:

```
service mgmt-vmware restart
```

Changing the service console firewall security level does not affect existing connections. For example, if the firewall is set to low security and a backup is running on a port you didn't explicitly open, raising the firewall setting to high does not terminate the backup. Rather, because the firewall is configured to pass packets for previously established connections, the backup completes, releases the connection, and no further connections are accepted for the port.

Opening and Closing Ports in the Service Console Firewall

You can open service console firewall ports when you install third-party devices, services, and agents. Before you open ports to support the item you are installing, see vendor specifications to determine the necessary ports.

If you close a port, active sessions of the service associated with the port are not automatically disconnected when you close the port. For example, if a backup is executing and you close the port for the backup agent, the backup continues until it completes and the agent releases the connection.

Perform the following procedures only if you are opening or closing ports for services or agents not specifically configurable through the VI Client. For information on configuring additional ports in VirtualCenter, see "Opening Firewall Ports for Supported Services and Management Agents" on page 188.

CAUTION VMware supports opening and closing firewall ports only through the VI Client or the `esxcfg-firewall` command, as described below. Using any other methods or scripts to open and close firewall ports can lead to unexpected behavior.

To open a specific port in the service console firewall

1 Log in to the service console and acquire root privileges.

2 Execute the following command:

```
esxcfg-firewall --openPort <port_number>,tcp|udp,in|out,<port_name>
```

where:

- `port_number` is the vendor-specified port number.
- `tcp|udp` is the protocol. Select `tcp` for TCP traffic or `udp` for UDP traffic.
- `in|out` is the traffic direction. Select `in` to open the port for inbound traffic or `out` to open it for outbound traffic.
- `port_name` is a descriptive name. The name does not need to be unique, but it should be meaningful to help identify the service or agent using the port.

  For example:

  ```
  esxcfg-firewall --openPort 6380,tcp,in,Navisphere
  ```

3 Execute the following command to restart the `vmware-hostd` process:

```
service mgmt-vmware restart
```

To close a specific port in the service console firewall

1 Log in to the service console and acquire root privileges.

2 Execute the following command:

```
esxcfg-firewall -closePort <port_number>,tcp|udp,in|out,<port_name>
```

The `port_name` argument is optional for `-closePort`.

For example:

```
esxcfg-firewall --closePort 6380,tcp,in
```

3 Execute the following command to restart the `vmware-hostd` process:

```
service mgmt-vmware restart
```

You can use the `-closePort` option to close only those ports that you opened with the `-openPort` option. If you used a different method to open the port, use an equivalent method to close it. For example, you can close the SSH port (22) only by disabling the SSH Server incoming connection and SSH Client outgoing connection in the VI Client. For information on opening and closing ports through the VI Client, see "Opening Firewall Ports for Supported Services and Management Agents" on page 188.

Password Restrictions

The ease with which an attacker can log in to an ESX Server 3 host depends on their ability to find a legitimate user name and password combination. A malicious user can obtain a password in a number of ways. For example, an attacker can sniff insecure network traffic, such as Telnet or FTP transmissions, for successful login attempts.

Another common method is to crack the password by running a password generator. Passwords generators are useful for mounting various kinds of password attacks, including brute force attacks, in which the generator tries every character combination up to a certain password length, and dictionary attacks, in which the generator tries real words and simple mutations of real words.

Implementing restrictions that govern the length, character sets, and duration of passwords can make attacks that a password generator initiates far more difficult. The longer and more complex the password, the harder it is for an attacker to discover. The more often users have to change passwords, the more difficult it is to find a password that works repeatably.

NOTE Always consider the human factor when you decide how to implement password restrictions. If you make passwords too hard to remember or enforce frequent password changes, your users might be inclined to write down their passwords, thus eliminating any benefit.

To help protect your password database from misuse, password shadowing is enabled for ESX Server 3 so that password hashes are hidden from access. Also, ESX Server 3 uses MD5 password hashes, which provide stronger password security and let you set minimum length requirements to more than eight characters.

ESX Server 3 provides password controls on two levels to help you enforce password policies for your users and limit the risk of password cracking:

- **Password aging** – These controls govern how long a user password can be active before the user is required to change it. They help ensure that passwords change often enough so that if an attacker obtains a password through sniffing or social engineering, they cannot keep accessing ESX Server 3 indefinitely.
- **Password complexity** – These controls ensure that the users select passwords that are hard for password generators to determine.

Password Aging

To ensure that passwords don't stay active for long periods, ESX Server 3 imposes the following password aging restrictions for user logins by default:

- **Maximum days** – (SEE UPDATE) The number of days that a user can keep a password before it needs to be changed. The default setting for ESX Server 3 is 90 days. By default, the `root` account and other service accounts are exempt from the 90 day expiration.
- **Minimum days** – The minimum number of days between password changes. The default setting is 0, meaning that the users can change their passwords any time.
- **Warning time** – The number of days in advance of password expiration that ESX Server 3 observes when issuing a password change reminder. The default setting is seven days. Warnings are only displayed when logging directly onto the service console or when using SSH.

You can tighten or loosen any of these settings by executing the `esxcfg-auth` command options. To override the default password aging settings for an individual user, use the `chage` command.

To change default password aging restrictions for ESX Server 3

1 Log in to the service console and acquire root privileges.

2 Execute one or more of the following commands as applicable.

- To change the maximum number of days a user can keep a password:

  ```
  esxcfg-auth --passmaxdays=<number_of_days>
  ```

  where <number_of_days> is the maximum number of days before password expiration.

- To change the minimum number of days between password changes:

  ```
  esxcfg-auth --passmindays=<number_of_days>
  ```

  where <number_of_days> is the minimum number of days between password changes.

- To change the warning time before a password change:

  ```
  esxcfg-auth --passwarnage=<number_of_days>
  ```

  where <number_of_days> is the number of days of advanced warning a user receives before a password change is due.

To override default password aging restrictions for individual users or groups

1 Log in to the service console and acquire root privileges.

2 Execute one or more of the following commands as applicable:

- To specify a new maximum days value:

  ```
  chage -M <number_of_days> <username>
  ```

- To specify a new minimum days value:

  ```
  chage -m <number_of_days> <username>
  ```

- To specify a new warning time value:

  ```
  chage -W <number_of_days> <username>
  ```

  To learn about other `chage` options, use the `man chage` command.

Password Complexity

By default, ESX Server 3 uses the `pam_cracklib.so` plug-in to set the rules that users must observe when creating passwords and to check password strength during the creation process.

The `pam_cracklib.so` plug-in lets you determine the basic standards that all passwords must meet. By default, ESX Server 3 imposes no restrictions on the root password. However, when nonroot users attempt to change their passwords, the passwords they choose must meet the basic standards that `pam_cracklib.so` sets. In addition, nonroot users can make only a certain number of password change attempts before `pam_cracklib.so` begins issuing messages and eventually closes the password change page. The ESX Server 3 defaults for `pam_cracklib.so` password standards and retry restrictions are as follows:

- **Minimum length** – The `pam_cracklib.so` minimum length parameter for ESX Server 3 systems is set to nine. This means that the user must enter at least eight characters if they use only one character class (lowercase, uppercase, digit, or other).

  The password length algorithm allows shorter passwords if the user enters a mix of character classes. To calculate the actual character length a user needs to enter to form a valid password for a given minimum length setting, apply the password length algorithm as follows:

  M – CC = E

  where:

  - M is the minimum length parameter.
  - CC is the number of character classes the user includes in the password.
  - E is the number of characters the user must enter.

Table 12-2 shows how the algorithm works, assuming the user enters at least one lowercase character as part of the password. The `pam_cracklib.so` plug-in does not allow passwords of fewer than six characters. Thus, while the mathematically accurate character requirement for a four character-class password is five characters, the effective requirement is six.

Table 12-2. Password Complexity Algorithm Results

# of Characters for a Valid Password	Character Types in the Password Attempt			
	Lowercase Characters	Uppercase Characters	Digits	Other Characters
8	yes			
7	yes	yes		
	yes		yes	
	yes			yes
6	yes	yes	yes	
	yes	yes		yes
	yes		yes	yes
5	yes	yes	yes	yes

- **Retries** – The `pam_cracklib.so` retries parameter for ESX Server 3 systems is set to three. If the user does not enter a strong enough password in three attempts, `pam_cracklib.so` closes the password change dialog box. The user must open a new password change session to try again.

The `pam_cracklib.so` plug-in checks all password change attempts to ensure that passwords meet the following strength criteria:

- The new password must not be a palindrome—a password where the characters mirror each other around a central letter, as in radar or civic.
- The new password must not be the reverse of the old password.
- The new password must not be a rotation—a version of the old password in which one or more characters have been rotated to the front or back of the password string.
- The new password must differ from the old password by more than a change of case.
- The new password must differ from the old password by more than a few characters.

- The new password must not have been used in the past. The `pam_cracklib.so` plug-in applies this criterion only if you have configured a password reuse rule.

  By default, ESX Server 3 does not enforce any password reuse rules, so ordinarily the `pam_cracklib.so` plug-in never rejects a password change attempt on these grounds. However, you can configure a reuse rule to ensure that your users do not alternate between a few passwords.

  If you configure a reuse rule, old passwords are stored in a file that the `pam_cracklib.so` plug-in references during each password change attempt. The reuse rules determine the number of old passwords that ESX Server 3 retains. When a user creates enough passwords to reach the value specified in the reuse rule, old passwords are removed from the file in age order. To learn how to configure a reuse rule, see "To configure a password reuse rule" on page 246.

- The new password must be long enough and complex enough. You configure these requirements by changing the `pam_cracklib.so` complexity parameters with the `esxcfg-auth` command, which lets you set the number of retries, the minimum password length, and a variety of character credits. Character credits let the user enter shorter passwords if they include more character types in the password. To learn how to configure password length and complexity, see "To change default password complexity for the pam_cracklib.so plug-in" on page 247.

For more information on the `pam_cracklib.so` plug-in, see your Linux documentation.

NOTE The `pam_cracklib.so` plug-in used in Linux provides more parameters than the parameters supported for ESX Server 3. You cannot specify these additional parameters in `esxcfg-auth`.

To configure a password reuse rule

1 Log in to the service console and acquire root privileges.

2 Change directories by entering `cd /etc/pam.d/` at the command prompt.

3 Use a text editor to open the `system-auth` file.

4 Locate the line that starts with:

```
password    sufficient    /lib/security/$ISA/pam_unix.so
```

5 Add the following parameters to the end of the line:

```
remember=X
```

where `X` is the number of old passwords to store for each user. Use a space as the delimiter between `remember=X` and the preceding parameter.

6 Save your changes and close the file.

7 Change directories to `/etc/security/` and issue the following command to make a zero (0) length file with `opasswd` as the filename:

```
touch opasswd
```

8 Enter the following commands:

```
chmod 0600 opasswd
chown root:root /etc/security/opasswd
```

To change default password complexity for the pam_cracklib.so plug-in

1 Log in to the service console and acquire root privileges.

2 Enter the following command:

```
esxcfg-auth --usecrack=<retries> <minimum_length> <lc_credit> <uc_credit>
<d_credit> <oc_credit>
```

where:

- `retries` is the number of retries the user is allowed before ESX Server 3 locks them out of password change mode.
- `minimum_length` is the minimum number of characters a user must enter to make the password acceptable. This number is the total length before any length credits are applied.

  One length credit is always applied so, in effect, the password length is one character less than the `minimum_length` parameter you specify. Because the `pam_cracklib.so` plug-in does not accept passwords of fewer than six characters, calculate the `minimum_length` parameter so that users cannot drop the password length below six as a result of subtracting the length credits.
- (SEE UPDATE) `lc_credit` is the number by which the `minimum_length` parameter is reduced if the user includes at least one lowercase character in the password.
- `uc_credit` is the number by which the `minimum_length` parameter is reduced if the user includes at least one uppercase character.
- `d_credit` is the number by which the `minimum_length` parameter is reduced if the user includes at least one digit.
- `oc_credit` is the number by which the `minimum_length` parameter is reduced if the user includes at least one special character, such as an underscore or dash.

Enter character credit parameters as a positive number or as zero (0) if you do not want the plug-in to give the user credit for including this character class. Character credits are additive. The more different types of characters the user enters, the fewer characters are required to form a valid password. For example, you issue the following command:

```
esxcfg-auth --usecrack=3 11 1 1 1 2
```

With this setting in effect, a user creating a password that contains lowercase characters and one underscore needs eight characters to create a valid password. If the user decides to include all types of characters (lowercase alphabetical, uppercase alphabetical, numeric, and special), they would need only six characters.

Changing the Password Plug-In

The `pam_cracklib.so` plug-in provides sufficient password strength enforcement for most environments. However, if the `pam_cracklib.so` plug-in is not stringent enough for your needs, you can use the `pam_passwdqc.so` plug-in instead. You change the plug-in through the `esxcfg-auth` command.

The `pam_passwdqc.so` plug-in tests for the same password characteristics as the `pam_cracklib.so` plug-in. However, it provides a greater number of options for fine-tuning password strength and performs password strength tests for all users, including the root user. The `pam_passwdqc.so` plug-in is also somewhat more difficult to use than the `pam_cracklib.so` plug-in. For more information on this plug-in, see your Linux documentation.

NOTE The `pam_passwdqc.so` plug-in used in Linux provides more parameters than the parameters supported for ESX Server 3. You cannot specify these additional parameters in `esxcfg-auth`.

To switch to the pam_passwdqc.so plug-in

1 Log in to the service console and acquire root privileges.

2 Enter the following command:

```
esxcfg-auth --usepamqc=<N0> <N1> <N2> <N3> <N4> <match>
```

where:

- `N0` is the number of characters required for a password that uses characters from only one character class.
- `N1` is the number of characters required for a password that uses characters from two character classes.

- `N2` is used for passphrases. ESX Server 3 requires three words for a passphrase.
- `N3` is the number of characters required for a password that uses characters from three character classes.
- `N4` is the number of characters required for a password that uses characters from all four character classes.
- `match` is the number of characters allowed in a string that is reused from the old password. If the `pam_passwdqc.so` plug-in finds a reused string of this length or longer, it disqualifies the string from the strength test and uses only the remaining characters.

Setting any of these options to `–1` directs the `pam_passwdqc.so` plug-in to ignore the requirement. Setting any of these options to `disabled` directs the `pam_passwdqc.so` plug-in to disqualify passwords with the associated characteristic. The values used must be in descending order except for `–1` and `disabled`.

For example, you issue the following command:

```
esxcfg-auth --usepamqc=disabled 18 -1 12 8
```

With this setting in effect, a user creating a password would never be able to set passwords that contain only one character class. The user needs to use at least 18 characters for a password with a two-character class, 12 characters for a three-character class password, and eight characters for four-character class passwords. Attempts to create passphrases are ignored.

Cipher Strength

Transmitting data over insecure connections presents a security risk because malicious users might be able to scan data as it travels through the network. As a safeguard, network components commonly encrypt the data so that it can't be easily read. To encrypt data, the sending component, such as a gateway or redirector, applies algorithms, or *ciphers*, to alter the data before transmitting it. The receiving component uses a key to decrypt the data, returning it to its original form.

Several different ciphers are currently in use, and the level of security that each provides is different. One measure of a cipher's ability to protect data is its *cipher strength*—the number of bits in the encryption key. The larger the number, the more secure the cipher.

To ensure the protection of the data transmitted to and from external network connections, ESX Server 3 uses one of the strongest block ciphers available—256-bit AES block encryption. ESX Server 3 also uses 1024-bit RSA for key exchange. These encryption algorithms are the default for the following connections:

- VI Client connections to the VirtualCenter Server and to the ESX Server 3 host through the service console.
- VI Web Access connections to the ESX Server 3 host through the service console.

  NOTE Because VI Web Access cipher usage is determined by the Web browser you are using, this management tool might use other ciphers.

- SDK connections to the VirtualCenter Server and to the ESX Server 3.
- Service console connections to virtual machines through the VMkernel.
- SSH connections to the ESX Server 3 host through the service console. For more information, see "SSH Security" on page 253.

setuid and setgid Applications

`setuid` is a flag that allows an application to temporarily change the permissions of the user running the application by setting the effective user ID to the program owner's user ID. `setgid` is a flag that allows an application to temporarily change the permissions of the group running the application by setting the effective group ID to the program owner's group ID.

During ESX Server 3 installation, several applications that include the `setuid` and `setgid` flags are installed by default. These applications are initiated by or through the service console. Some of them provide facilities required for correct operation of the ESX Server 3 host. Others are optional, but they can make maintaining and troubleshooting the ESX Server 3 host and the network easier.

Default setuid Applications

Table 12-3 lists the default `setuid` applications and indicates whether the application is required or optional.

Table 12-3. Default setuid Applications

Application	Purpose and Path	Required or Optional
`crontab`	Lets individual users add cron jobs. **Path:** `/usr/bin/crontab`	Optional
`pam_timestamp_check`	Supports password authentication. **Path:** `/sbin/pam_timestamp_check`	Required
`passwd`	Supports password authentication. **Path:** `/usr/bin/passwd`	Required
`ping`	Sends and listens for control packets on the network interface. Useful for debugging networks. **Path:** `/bin/ping`	Optional
`pwdb_chkpwd`	Supports password authentication. **Path:** `/sbin/pwdb_chkpwd`	Required
`ssh-keysign`	Performs host-based authentication for SSH. **Path:** `/usr/libexec/openssh/ssh-keysign`	Required if you use host-based authentication. Otherwise optional.
`su`	Lets a general user become the root user by changing users. **Path:** `/bin/su`	Required
`sudo`	Lets a general user act as the root user only for specific operations. **Path:** `/usr/bin/sudo`	Optional
`unix_chkpwd`	Supports password authentication. **Path:** `/sbin/unix_chkpwd`	Required
`vmkload_app`	Performs tasks required to run virtual machines. This application is installed in two locations: one for standard use and one for debugging. **Path for standard use:** `/usr/lib/vmware/bin/vmkload_app` **Path for debugging:** `/usr/lib/vmware/bin-debug/vmkload_app`	Required in both paths

Table 12-3. Default setuid Applications (Continued)

Application	Purpose and Path	Required or Optional
`vmware-authd`	Authenticates users for use of services specific to VMware. **Path:** `/usr/sbin/vmware-authd`	Required
`vmware-vmx`	Performs tasks required to run virtual machines. This application is installed in two locations: one for standard use and one for debugging. **Path for standard use:** `/usr/lib/vmware/bin/vmware-vmx` **Path for debugging:** `/usr/lib/vmware/bin-debug/vmware-vmk`	Required in both paths

Disabling any of the required applications results in problems with ESX Server 3 authentication and virtual machine operation, but you can disable any optional application.

To disable an optional setuid application

1 Log in to the service console and acquire root privileges.

2 Execute the following command:

```
chmod a-s <path_to_executable>
```

Default setgid Applications

Two applications that include the `setgid` flag are installed by default. Table 12-4 lists the default `setgid` applications and indicates whether the application is required or optional.

Table 12-4. Default `setgid` Applications

Application	Purpose and Path	Required or Optional
`wall`	Alerts all terminals that an action is about to occur. This application is called by `shutdown` and other commands. **Path:** `/usr/bin/wall`	Optional
`lockfile`	Performs locking for the Dell OM management agent. **Path:** `/usr/bin/lockfile`	Required for Dell OM but optional otherwise

Disabling a required application results in problems with ESX Server 3 authentication and virtual machine operation, but you can disable any optional application.

To disable an optional setgid application

1 Log in to the service console and acquire root privileges.

2 Execute the following command:

```
chmod a-g <path_to_executable>
```

SSH Security

SSH is a commonly used UNIX and Linux command shell that lets you remotely log in to the service console and perform certain management and configuration tasks for ESX Server 3. SSH is used for secure logins and data transfers because it offers stronger protection than other command shells. In this ESX Server 3 release, the SSH configuration is enhanced to provide a higher security level. Key features of this enhancement include:

- **Version 1 SSH protocol disabled** – VMware no longer supports Version 1 SSH protocol and uses Version 2 protocol exclusively. Version 2 eliminates certain security issues present in Version 1 and provides you with a safer communications interface to the service console.
- **Improved cipher strength** – SSH now supports only 256-bit and 128-bit AES ciphers for your connections.
- **Limits on remote logins as root** – You can no longer remotely log in as root. Instead, you log in as an identifiable user and either use the `sudo` command to execute specific operations that require root privileges or enter the `su` command to become the root user.

NOTE The `sudo` command provides security benefits in that it limits root activities and helps you check for possible misuse of root privileges by generating an audit trail of any root activities that the user performs.

These settings are designed to provide solid protection for the data you transmit to the service console through SSH. If this configuration is too rigid for your needs, you can lower security parameters.

To change the default SSH configuration

1 Log in to the service console and acquire root privileges.

2 Change directories by entering **`cd /etc/ssh`** at the command prompt.

3 Use a text editor to perform any or all of following actions, as appropriate.

- To allow remote root login, change the setting to `yes` in the following line in the `sshd_config` file:

  ```
  PermitRootLogin no
  ```

- To revert to the default SSH protocol (Version 1 and 2), comment out the following line in the `sshd_config` file:

  ```
  Protocol 2
  ```

- To revert to the 3DES cipher and other ciphers, comment out the following line in the `sshd_config` file:

  ```
  Ciphers aes256-cbc,aes128-cbc
  ```

- To disable Secure FTP (SFTP) on SSH, comment out the following line in the `sshd_config` file:

  ```
  Subsystem ftp /usr/libexec/openssh/sftp-server
  ```

4 Save your changes and close the file.

5 Execute the following command to restart the SSHD service:

```
service sshd restart
```

Security Patches and Security Vulnerability Scanning Software

If a fix for a particular LINUX-supported software package that VMware provides as a service console component becomes available—for example, a service, facility, or protocol—VMware provides an RPM Package Manager (RPM) package that you use to update the software package on ESX Server 3. Although these fixes might be available from other sources, always use RPMs that VMware generates instead of using third-party RPMs.

When providing patches for a software package, the VMware policy is to backport the fix to a version of the software known to be stable. This approach reduces the chance of introducing new problems and instability in the software. Because the patch is added to an existing version of the software, the version number of the software stays the same, but a patch number is added as a suffix.

Certain security scanners such as Nessus check the version number but not the patch suffix as they search for security holes. As a result, these scanners can falsely report that software is down-level and doesn't include the most recent security patches even though it does. This problem is common to the industry and not specific to VMware.

NOTE Some security scanners can handle this situation correctly, but they typically lag by a version or more. For example, the version of Nessus released after a Red Hat patch often doesn't report these false positives.

The following is an example of how this problem occurs:

1 You initially install ESX Server 3 with OpenSSL version 0.9.7a (where 0.9.7a is the original version with no patches).

2 OpenSSL releases a patch that fixes a security hole in version 0.9.7. This version is called 0.9.7x.

3 VMware backports the OpenSSL 0.9.7x fix to the original version, updates the patch number, and creates an RPM. The OpenSSL version in the RPM is 0.9.7a-1, indicating that the original version (0.9.7a) now contains patch 1.

4 You install the RPM.

5 The security scanner fails to note the -1 suffix and erroneously reports that security for OpenSSL isn't up-to-date.

If your scanner reports that security for a package is down-level, perform the following checks:

- Look at the patch suffix to determine whether you need to get an update.
- Read the VMware RPM documentation for information on the patch contents.
- Use the following command to look for Common Vulnerabilities and Exposures (CVE) number from the security alert in the RPM change log:

  ```
  rpm– q ––changelog openssl | grep <CVE_number>
  ```

  If the CVE number is there, the specified package addresses that vulnerability.

Security Deployments and Recommendations

13

The chapter focuses on giving you a better idea of how to secure your ESX Server 3 in particular environments by presenting a series of ESX Server 3 deployment scenarios that you can consider as you plan some of the security features of your own deployment. It also makes some basic security recommendations you can consider when creating and configuring virtual machines.

This chapter covers the following topics:

- "Security Approaches for Common ESX Server 3 Deployments" on page 257
- "Virtual Machine Recommendations" on page 263

Security Approaches for Common ESX Server 3 Deployments

The complexity of ESX Server 3 deployments can vary significantly depending on the size of your company, the way that data and resources need to be shared with the outside world, whether there are multiple datacenters or just one, and so forth.

Inherent in the following deployments are policies for user access, resource sharing, and security level. By comparing the deployments, you can get a sense of the issues you face in planning security for your own ESX Server 3 deployment.

Single Customer Deployment

In this deployment, the ESX Server 3 hosts are owned and maintained within a single corporation and single datacenter. No ESX Server 3 resources are shared with outside users. One site administrator maintains the ESX Server 3 hosts, and these hosts run a number of virtual machines.

The deployment does not allow customer administrators, and the site administrator is solely responsible for maintaining the various virtual machines. The corporation staffs a set of system administrators who do not have accounts on the ESX Server 3 host and cannot access any of the ESX Server 3 tools such as VirtualCenter or command line shells for the host. These system administrators have access to virtual machines through the virtual machine console so that they can load software and perform other maintenance tasks inside the virtual machines.

Table 13-1 shows how you might handle sharing for the components you use and configure for the ESX Server 3 host.

Table 13-1. Sharing for Components in a Single Customer Deployment

Function	Configuration	Comments
Service console shares the same physical network as the virtual machines?	No	Isolate the service console by configuring it on its own physical network.
Service console shares the same VLAN as the virtual machines?	No	Isolate the service console by configuring it on its own VLAN. No virtual machine or other system facility such as VMotion should use this VLAN.
Virtual machines share the same physical network?	Yes	Configure your virtual machines on the same physical network.
Network adapter sharing?	Partial	Isolate the service console by configuring it on its own virtual switch and virtual network adapter. No virtual machine or other system facility should use this switch or adapter. However, you can configure your virtual machines on the same virtual switch and network adapter.
VMFS sharing?	Yes	All `.vmdk` files should reside in the same VMFS partition.
Security level	High	Open ports for needed services like FTP on an individual basis. See "Service Console Firewall Configuration" on page 237 for information on security levels.
Virtual machine memory overcommitment?	Yes	Configure the total memory for the virtual machines as greater than the total physical memory.

Table 13-2 shows how you might set up user accounts for the ESX Server 3 host.

Table 13-2. User Account Setup in a Single Customer Deployment

User Category	Total Number of Accounts
Site administrators	1
Customer administrators	0
System administrators	0
Business users	0

Table 13-3 shows the level of access for each user.

Table 13-3. User Access in a Single Customer Deployment

Access Level	Site Administrator	System Administrator
Root access?	Yes	No
Service console access through SSH?	Yes	No
VirtualCenter and VI Web Access?	Yes	No
Virtual machine creation and modification?	Yes	No
Virtual machine access through the console?	Yes	Yes

Multiple Customer Restricted Deployment

In this deployment, the ESX Server 3 hosts are in the same datacenter and are used to serve applications for multiple customers. The site administrator maintains the ESX Server 3 hosts, and these hosts run a number of virtual machines dedicated to the customers. Virtual machines that belong to the various customers can be on the same ESX Server 3 host, but the site administrator restricts resource sharing to prevent rogue interaction.

While there is only one site administrator, several customer administrators maintain the virtual machines assigned to their customers. This deployment also includes customer system administrators who do not have ESX Server 3 accounts but have access to the virtual machines through the virtual machine console so that they can load software and perform other maintenance tasks inside the virtual machines.

Table 13-4 shows how you might handle sharing for the components you use and configure for the ESX Server 3 host.

Table 13-4. Sharing for Components in a Multiple Customer Restricted Deployment

Function	Configuration	Comments
Service console shares the same physical network as the virtual machines?	No	Isolate the service console by configuring it on its own physical network.
Service console shares the same VLAN as the virtual machines?	No	Isolate the service console by configuring it on its own VLAN. No virtual machine or other system facility such as VMotion should use this VLAN.
Virtual machines share the same physical network?	Partial	Put the virtual machines for each customer on a different physical network. All physical networks are independent of each other.
Network adapter sharing?	Partial	Isolate the service console by configuring it on its own virtual switch and virtual network adapter. No virtual machine or other system facility should use this switch or adapter. You configure virtual machines for one customer so that they all share the same virtual switch and network adapter. However, they do not share the switch and adapter with any other customers.
VMFS sharing?	No	Each customer has their own VMFS partition and their virtual machine `.vmdk` files reside exclusively on that partition. The partition can span multiple LUNs.
Security level	High	Open ports for services like FTP as needed.
Virtual machine memory overcommitment?	Yes	Configure the total memory for the virtual machines as greater than the total physical memory.

Table 13-5 shows how you might set up user accounts for the ESX Server 3 host.

Table 13-5. User Account Setup in a Multiple Customer Restricted Deployment

User Category	Total Number of Accounts
Site administrators	1
Customer administrators	10
System administrators	0
Business users	0

Table 13-6 shows the level of access for each user.

Table 13-6. User Access in a Multiple Customer Restricted Deployment

Access Level	Site Administrator	Customer Administrator	System Administrator
Root access?	Yes	No	No
Service console access through SSH?	Yes	Yes	No
VirtualCenter and VI Web Access?	Yes	Yes	No
Virtual machine creation and modification?	Yes	Yes	No
Virtual machine access through the console?	Yes	Yes	Yes

Multiple Customer Open Deployment

In this deployment, the ESX Server 3 hosts are in the same datacenter and are used to serve applications for multiple customers. The site administrator maintains the ESX Server 3 hosts, and these hosts run a number of virtual machines dedicated to the customers. Virtual machines that belong to the various customers can be on the same ESX Server 3 host, but there are fewer restrictions on resource sharing.

While there is only one site administrator, several customer administrators maintain the virtual machines assigned to their customers. The deployment also includes customer system administrators who do not have ESX Server 3 accounts but have access to the virtual machines through the virtual machine console so that they can load software and perform other maintenance tasks inside the virtual machines. Lastly, a group of business users who do not have accounts can use virtual machines to run their applications.

Table 13-7 shows how you might handle sharing for the components you use and configure for the ESX Server 3 host.

Table 13-7. Sharing for Components in a Multiple Customer Open Deployment

Function	Configuration	Comments
Service console shares the same physical network as the virtual machines?	No	Isolate the service console by configuring it on its own physical network.
Service console shares the same VLAN as the virtual machines?	No	Isolate the service console by configuring it on its own VLAN. No virtual machine or other system facility such as VMotion should use this VLAN.
Virtual machines share the same physical network?	Yes	Configure your virtual machines on the same physical network.
Network adapter sharing?	Partial	Isolate the service console by configuring it on its own virtual switch and virtual network adapter. No virtual machine or other system facility should use this switch or adapter. You configure all virtual machines on the same virtual switch and network adapter.
VMFS sharing?	Yes	Virtual machines can share VMFS partitions and their virtual machine `.vmdk` files can reside on shared partitions. Virtual machines do not share `.vmdk` files.
Security level	High	Open ports for services like FTP as needed.
Virtual machine memory overcommitment?	Yes	Configure the total memory for the virtual machines as greater than the total physical memory.

Table 13-8 shows how you might set up user accounts for the ESX Server 3 host.

Table 13-8. User Account Setup in a Multiple Customer Open Deployment

User Category	Total Number of Accounts
Site administrators	1
Customer administrators	10
System administrators	0
Business users	0

Table 13-9 shows the level of access for each user.

Table 13-9. User Access in a Multiple Customer Open Deployment

Access Level	Site Administrator	Customer Administrator	System Administrator	Business User
Root access?	Yes	No	No	No
Service console access through SSH?	Yes	Yes	No	No
VirtualCenter and VI Web Access?	Yes	Yes	No	No
Virtual machine creation and modification?	Yes	Yes	No	No
Virtual machine access through the console?	Yes	Yes	Yes	Yes

Virtual Machine Recommendations

Consider the following safety precautions when evaluating virtual machine security and administering virtual machines.

Installing Antivirus Software

Because each virtual machine hosts a standard operating system, you should consider protecting it from viruses by installing antivirus software. Depending on how you are using the virtual machine, you might also want to install a software firewall.

NOTE Software firewalls and antivirus software can be virtualization-intensive. If you are confident that your virtual machines are in a fully trusted environment, you can balance the need for these two security measures against virtual machine performance.

Disabling Copy and Paste Operations Between the Guest Operating System and Remote Console

When VMware Tools runs on a virtual machine, you can copy and paste between the guest operating system and remote console. As soon as the console window gains focus, non-privileged users and processes running in the virtual machine can access the clipboard for the virtual machine console. If a user copies sensitive information to the clipboard before using the console, the user—perhaps unknowingly—exposes sensitive data to the virtual machine.

To prevent this problem, consider disabling copy and paste operations for the guest operating system.

To disable copy and paste operations between the guest operating system and remote console

1 Log on to the VI Client and select the virtual machine from the inventory panel.

The configuration page for this virtual machine appears with the **Summary** tab displayed.

2 Click **Edit Settings**.

3 Click **Options** > **Advanced** > **Configuration Parameters** to open the **Configuration Parameters** dialog box.

4 Click the **Add** button.

5 Type the following values in the **Name** field **Value** column.

Table 13-10. Configuration Parameter Settings

Name Field	Value Field
`isolation.tools.copy.disable`	`true`
`isolation.tools.paste.disable`	`true`
`isolation.tools.setGUIOptions.enable`	`false`

The result appears as follows.

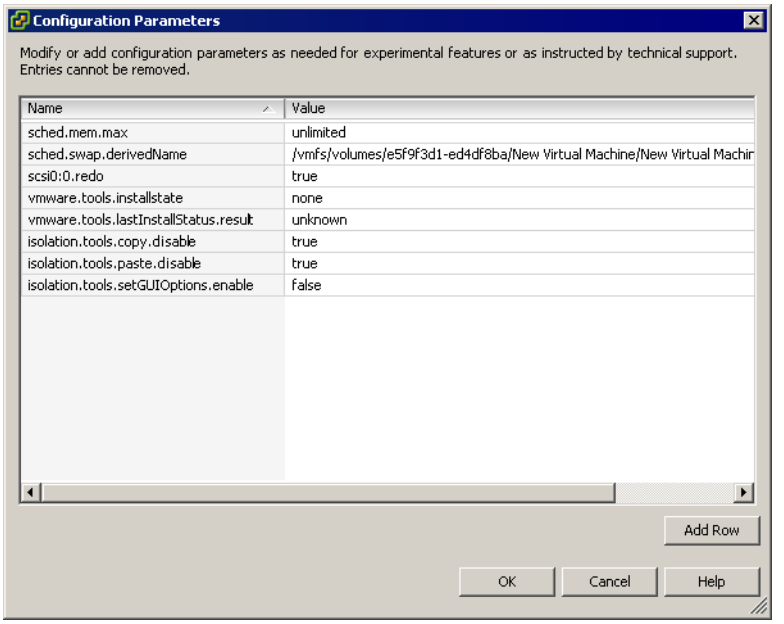

NOTE These options override any settings made in the guest operating system's VMware Tools control panel.

6 Click **OK** to close the **Configuration Parameters** dialog box and then click **OK** again to close the **Virtual Machine Properties** dialog box.

Removing Unnecessary Hardware Devices

Nonprivileged users and processes within virtual machines can connect or disconnect hardware devices, such as network adapters and CD-ROM drives. Attackers can use this capability to breach virtual machine security in several ways. For example, by default, an attacker with access to a virtual machine can:

- Connect a disconnected CD-ROM drive and access sensitive information on the media left in the drive.
- Disconnect a network adapter to isolate the virtual machine from its network, resulting in a denial of service.

As a general security precaution, use commands on the VI Client Configuration tab to remove any unneeded or unused hardware devices. While this measure tightens virtual machine security, it isn't a good solution in situations where you might need to bring a currently unused device back into service at a later time.

If you don't want to permanently remove a device, you can prevent a virtual machine user or process from connecting or disconnecting the device from within the guest operating system.

To prevent a virtual machine user or process from disconnecting devices

1 Log on to the VI Client and select the virtual machine from the inventory panel.

The configuration page for this virtual machine appears with the **Summary** tab displayed.

2 Click **Edit Settings**.

The **Virtual Machine Properties** dialog box appears.

3 Click **Options** > **General** and make a record of the path displayed in the **Virtual Machine Configuration File** field.

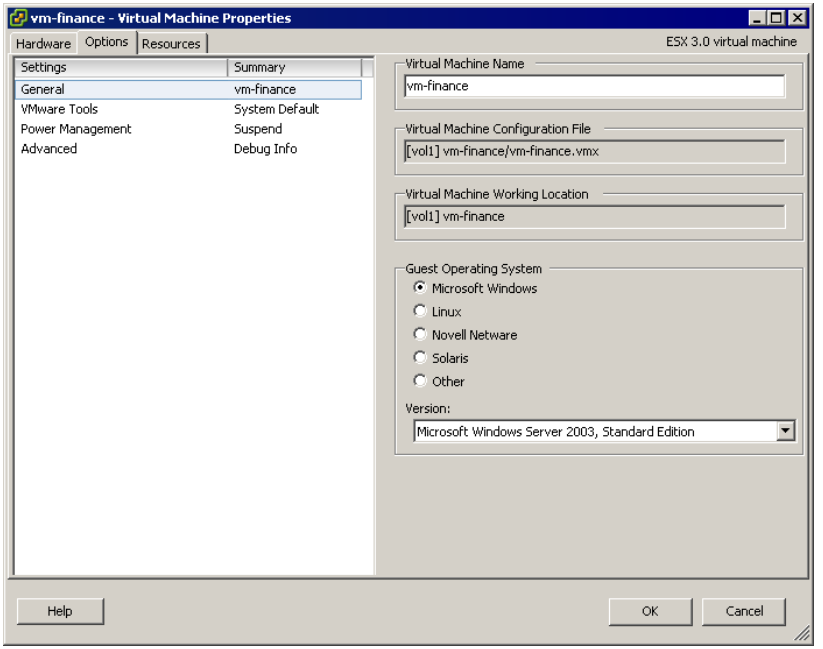

4 Log on to the service console and acquire root privileges.

5 Change directories to access the virtual machine configuration file whose path you recorded in Step 3.

Virtual machine configuration files are located in the `/vmfs/volumes/<datastore>` directory, where `<datastore>` is the name of the storage device on which the virtual machine files reside. For example, if the virtual machine configuration file you obtained from the **Virtual Machine Properties** dialog box is `[vol1]vm-finance/vm-finance.vmx`, you change directories as follows:

```
cd /vmfs/volumes/vol1/vm-finance/
```

6 Use nano or another text editor to add the following line to the `.vmx` file.

```
<device_name>.allowGuestConnectionControl = "false"
```

Where `<device_name>` is the name of the device you want to protect, for example, ethernet1.

NOTE By default, Ethernet 0 is configured to disallow device disconnection. The only reason you might need to change this is if a prior administrator set the `<device_name>.allowGuestConnectionControl` to `true`.

---

7 Save your changes and close the file.

8 Return to the VI Client and power off and power on the virtual machine. To do so, right-click the virtual machine in the inventory panel and click **Power Off** followed by **Power On**.

Limiting Guest Operating System Writes to Host Memory

The guest operating system processes send informational messages to the ESX Server 3 host through VMware Tools. These messages, known as `setinfo` messages, typically contain name-value pairs that define virtual machine characteristics or identifiers that the host stores—for example, `ipaddress=10.17.87.224`.

If the amount of data the host stored as a result of these messages was unlimited, an unrestricted data flow would provide an opportunity for an attacker to stage a DOS attack by writing software that mimics VMware Tools and filling the host's memory with arbitrary configuration data, thus consuming space needed by the virtual machines.

To prevent this problem, the configuration file containing these name-value pairs is limited to a size of one megabyte. One megabyte should be sufficient for most cases, but this value can be changed, as required. You might increase this value if large amounts of custom information are being stored in the configuration file.

To modify the GuestInfo memory limit, set the tools.setInfo.sizeLimit attribute of the.vmx file. The default limit is one megabyte, and this is applied, even if the `sizeLimit` attribute does not exist.

To modify guest operating system variable memory limit

1 Log on to the VI Client and select the virtual machine from the inventory panel.

The configuration page for this virtual machine appears with the **Summary** tab displayed.

2 Click **Edit Settings**.

3 Click **Options** > **Advanced** > **Configuration Parameters** to open the **Configuration Parameters** dialog box.

4 If the size limit attribute is not present, click **Add Row** and type the following:

- **Name field** – **`tools.setInfo.sizeLimit`**
- **Value field** – **`<Number of Bytes>`**

If the size limit attribute exists, modify it to reflect the desired limits.

A configuration that limits the GuestInfo size to 1048576 bytes (one MB) would appear as follows:

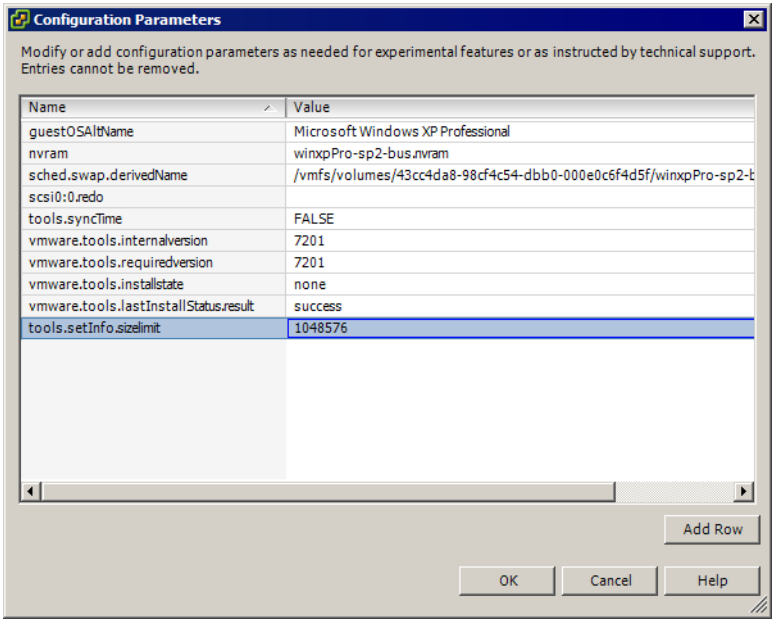

5 Click **OK** to close the **Configuration Parameters** dialog box, and then click **OK** again to close the **Virtual Machine Properties** dialog box.

You may also elect to entirely prevent guests from writing any name-value pairs to the configuration file. This is appropriate when guest operating systems must be prevented from modifying configuration settings.

To prevent the guest operating system processes from sending configuration messages to the host

1 Log on to the VI Client and select the virtual machine from the inventory panel.

The configuration page for this virtual machine appears with the **Summary** tab displayed.

2 Click **Edit Settings**.

3 Click **Options** > **Advanced** > **Configuration Parameters** to open the **Configuration Parameters** dialog box.

4 Click the **Add** button and type the following:

- **Name field – `isolation.tools.setinfo.disable`**
- **Value field – `true`**

The result appears as follows.

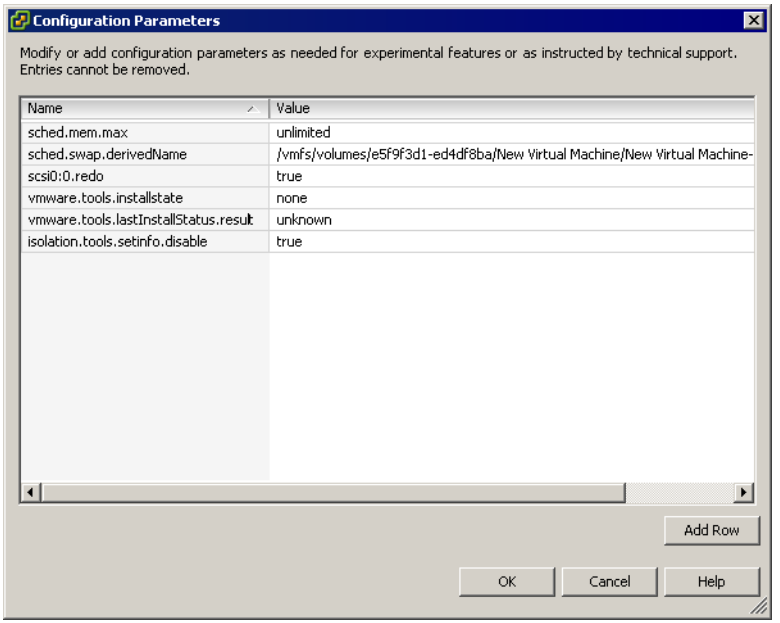

5 Click **OK** to close the **Configuration Parameters** dialog box and then click **OK** again to close the **Virtual Machine Properties** dialog box.

Configuring Logging Levels for the Guest Operating System

Virtual machines can write troubleshooting information into a virtual machine log file stored on the VMFS volume. Virtual machine users and processes can abuse logging either on purpose or inadvertently so that large amounts of data flood the log file. Over time, the log file can consume enough file system space to cause a denial of service.

To prevent this problem, consider modifying logging settings for virtual machine guest operating systems. These settings can limit the total size and number of log files. Normally a new log file is created each time a host is rebooted, so the file can grow to be quite large, but you can ensure new log file creation happens more frequently by limiting the maximum size of the log files. If you want to restrict the total size of logging

data, VMware recommends saving 10 log files, each one limited to 100KB. These values are small enough that the log files should not consume an undue amount of disk space on the host, yet the extent of data stored is large enough that it should capture sufficient information to debug most problems that might occur.

Each time an entry is written to the log, the size of the log is checked, and if it is over the limit, the next entry is written to a new log. If there are already the maximum number of log files, when a new one is created, the oldest one is deleted. A denial of service attack that avoids these limits could be attempted by writing an enormous log entry, but each log entry is limited in size to 4 KB, so no log files are ever more than 4 KB larger than the configured limit.

To limit log file numbers and sizes

1 Log on to the VI Client and select the virtual machine from the inventory panel.

   The configuration page for this virtual machine appears with the **Summary** tab displayed.

2 Click **Edit** Settings.

   The Virtual Machine Properties dialog box appears.

3 Click **Options** > **General** and make a record of the path displayed in the **Virtual Machine Configuration File** field.

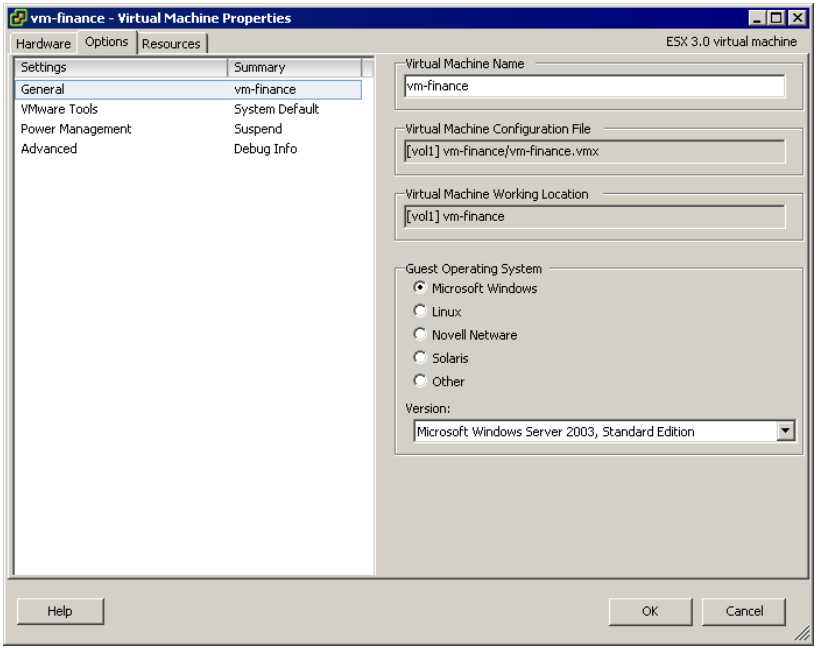

4 Log on to the service console and acquire root privileges.

5 Change directories to access the virtual machine configuration file whose path you recorded in Step 3.

Virtual machine configuration files are located in the `/vmfs/volumes/<datastore>` directory, where `<datastore>` is the name of the storage device on which the virtual machine files reside. For example, if the virtual machine configuration file you obtained from the **Virtual Machine Properties** dialog box is `[vol1]vm-finance/vm-finance.vmx`, you change directories as follows:

```
cd /vmfs/volumes/vol1/vm-finance/
```

6 To limit the log size, use nano or another text editor to add or edit the following line to the `.vmx` file.

```
log.rotateSize=<maximum size>
```

Where `<maximum size>` is the maximum file size in bytes. For example, to limit the size to around 100 KB, you could enter **`100000`**.

7 To keep a limited number of log files, use nano or another text editor to add or edit the following line to the `.vmx` file.

```
log.keepOld=<number of files to keep>
```

Where `<number of files to keep>` is the number of files the server will keep. For example, to keep 10 log files and then begin deleting the oldest ones as new ones are created, you could enter **`10`**.

It is also possible to stop logging altogether. In making this decision, be aware that you might not be able to gather adequate logs to allow troubleshooting. Further, VMware does not offer technical support for virtual machine problems if logging has been disabled.

To disable logging for the guest operating system

1 Log on to the VI Client and select the virtual machine from the inventory panel.

The configuration page for this virtual machine appears with the **Summary** tab displayed.

2 Click **Edit Settings**.

3 Click **Options** > **Advanced** > **General**.

4 Deselect the **Enable logging** check box.

The result appears as follows.

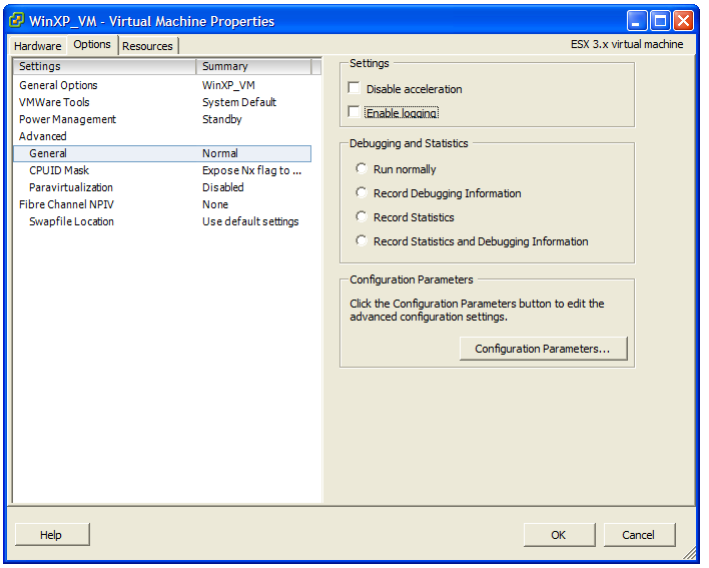

5 Click **OK** to close the **Virtual Machine Properties** dialog box.

Appendixes

ESX Server 3 Technical Support Commands

A

This appendix lists the service console commands used to configure ESX Server 3. Most of these commands are reserved for Technical Support use and are included for your reference only. In a few cases, however, these commands provide the only means of performing a configuration task for the ESX Server 3 host. Also, if you lose your connection to the host, executing certain of these commands through the command-line interface may be your only recourse—for example, if networking becomes nonfunctional and VI Client access is therefore unavailable.

NOTE If you use the commands in this appendix, you must execute the `service mgmt-vmware restart` command to restart the `vmware-hostd` process and alert the VI Client and other management tools that the configuration has changed. In general, avoid executing the commands in this appendix if the host is currently under the VI Client or VirtualCenter Server management.

The VI Client graphical user interface provides the preferred means of performing the configuration tasks described in this appendix. You can use this appendix to learn which VI Client commands to use in place of the service console commands. This appendix provides a summary of the actions you take in VI Client but doesn't give complete instructions. For details on using commands and performing configuration tasks through VI Client, see the online help.

You can find additional information on a number of ESX Server 3 commands by logging on to the service console and using the `man <esxcfg_command_name>` command to display man pages.

Table A-1 lists the Technical Support commands provided for ESX Server 3, summarizes the purpose of each command, and provides a VI Client alternative. You can perform most of the VI Client actions listed in the table only after you have selected an ESX Server 3 host from the inventory panel and clicked the **Configuration** tab. These actions are preliminary to any procedure discussed below unless otherwise stated.

Table A-1. ESX Server 3 Technical Support Commands

Service Console Command	Command Purpose and VI Client Procedure
`esxcfg-advcfg`	Configures advanced options for ESX Server 3. To configure advanced options in VI Client, click **Advanced Settings**. When the **Advanced Settings** dialog box opens, use the list on the left to select the device type or activity you want to work with and then enter the appropriate settings.
`esxcfg-auth`	Configures authentication. You can use this command to switch between the `pam_cracklib.so` and `pam_passwdqc.so` plugins for password change rule enforcement. You also use this command to reset options for these two plugins. For more information, see "Password Complexity" on page 244. There is no means of configuring these functions in VI Client.
`esxcfg-boot`	Configures bootstrap settings. This command is used for the bootstrap process and is intended for VMware Technical Support use only. You should not issue this command unless instructed to do so by a VMware Technical Support representative. There is no means of configuring these functions in VI Client.
`esxcfg-dumppart`	Configures a diagnostic partition or searches for existing diagnostic partitions. When you install ESX Server 3, a diagnostic partition is created to store debugging information in the event of a system fault. You don't need to create this partition manually unless you determine that there is no diagnostic partition for the host. You can perform the following management activities for diagnostic partitions in VI Client: ■ **Determine whether there is a diagnostic partition** — Click **Storage**>**Add** and check the first page of the **Add Storage** Wizard to see whether it includes the **Diagnostic** option. If **Diagnostic** is not one of the options, ESX Server 3 already has a diagnostic partition. ■ **Configure a diagnostic partition** — Click **Storage**>**Add**>**Diagnostic** and step through the wizard.

Table A-1. ESX Server 3 Technical Support Commands (Continued)

Service Console Command	Command Purpose and VI Client Procedure
`esxcfg-firewall`	Configures the service console firewall ports. To configure firewall ports for supported services and agents in VI Client, you select the Internet services that will be allowed to access the ESX Server 3 host. Click **Security Profile**>**Firewall**>**Properties** and use the **Firewall Properties** dialog box to add services. For details on adding supported services and configuring firewalls, see "Opening Firewall Ports for Supported Services and Management Agents" on page 188. You cannot configure unsupported services through the VI Client. For these services, use the `esxcfg-firewall` command as described in "Service Console Firewall Configuration" on page 237.
`esxcfg-info`	Prints information about the state of the service console, VMkernel, various subsystems in the virtual network, and storage resource hardware. VI Client doesn't provide a method for printing this information, but you can obtain much of it through different tabs and functions in the user interface. For example, you can check the status of your virtual machines by reviewing the information on the **Virtual Machines** tab.
`esxcfg-init`	Performs internal initialization routines. This command is used for the bootstrap process you should not use it under any circumstances. Using this command can cause problems for your ESX Server 3 host. There is no VI Client equivalent for this command.
`esxcfg-linuxnet`	Converts `vswif` to `eth` when booting ESX Server 3 into service-console-only mode rather than into ESX mode. This command is used for the bootstrap process and is intended for VMware Technical Support use only. You should not issue this command unless instructed to do so by a VMware Technical Support representative. There is no VI Client equivalent for this command.
`esxcfg-module`	Sets driver parameters and modifies which drivers are loaded during startup. This command is used for the bootstrap process and is intended for VMware Technical Support use only. You should not issue this command unless instructed to do so by a VMware Technical Support representative. There is no VI Client equivalent for this command.
`esxcfg-mpath`	Configures multipath settings for your Fibre Channel or iSCSI disks. To configure multipath settings for your storage in VI Client, click **Storage**. Select a datastore or mapped LUN and click **Properties**. When the **Properties** dialog box opens, select the desired extent if necessary. Then, click **Extent Device**>**Manage Paths** and use the **Manage Path** dialog box to configure the paths.

Table A-1. ESX Server 3 Technical Support Commands (Continued)

Service Console Command	Command Purpose and VI Client Procedure
`esxcfg-nas`	Manages NAS mounts. You use this command to add, delete, list, and change the attributes of NAS devices. To view NAS devices in VI Client, click **Storage** and scroll through the storage list. You can also perform the following activities from the **Storage** view: ■ **Display the attributes of a NAS device** – Click the device and review the information under **Details**. ■ **Add a NAS device** – Click **Add Storage**. ■ **Delete a NAS device** – Click **Remove**. ■ **Change the attributes of a NAS device** – Click the device and click **Details>Properties**. For complete instructions on how to create and configure NAS datastores, see "Configuring ESX Server 3 to Access NFS Volumes" on page 129.
`esxcfg-nics`	Prints a list of physical network adapters along with information on the driver, PCI device, and link state of each NIC. You can also use this command to control a physical network adapter's speed and duplexing. To view information on the physical network adapters for the host in VI Client, click **Network Adapters**. To change the speed and duplexing for a physical network adapter in the VI Client, click **Networking>Properties** for any of the virtual switches associated with the physical network adapter. In the **Properties** dialog box, click **Network Adapters>Edit** and select the speed and duplex combination. For complete instructions on how to change the speed and duplexing, see "To configure the uplink network adapter by changing its speed" on page 43.
`esxcfg-resgrp`	Restores resource group settings and lets you perform basic resource group management. Select a resource pool from the inventory panel and click **Edit Settings** on the **Summary** tab to change the resource group settings.
`esxcfg-route`	Sets or retrieves the default VMkernel gateway route and adds, removes, or lists static routes. To view the default VMkernel gateway route in VI Client, click **DNS and Routing**. To change the default routing, click **Properties** and update the information in both tabs of the **DNS and Routing Configuration** dialog box.

Table A-1. ESX Server 3 Technical Support Commands (Continued)

Service Console Command	Command Purpose and VI Client Procedure
`esxcfg-swiscsi`	Configures your software iSCSI software adapter. To configure your software iSCSI system in VI Client, click **Storage Adapters**, select the iSCSI adapter you want to configure, and click **Properties**. Use the **iSCSI Initiator Properties** dialog box to configure the adapter. For complete instructions on how to create and configure iSCSI datastores, see "iSCSI Storage" on page 110.
`esxcfg-upgrade`	Upgrades ESX Server from ESX Server 2.*x* to ESX Server 3. This command is not for general use. You complete the following three tasks when upgrading from 2.*x* to 3.*x*. Some of these can be performed in VI Client: ■ **Upgrade the host** — You upgrade the binaries, converting from ESX Server 2.*x* to ESX Server 3. You cannot perform this step from VI Client. For information on performing this upgrade, see the *Installation and Upgrade Guide*. ■ **Upgrade the file system** — To upgrade VMFS-2 to VMFS-3, suspend or power off your virtual machines and then click **Inventory>Host>Enter Maintenance Mode**. Click **Storage**, select a storage device, and click **Upgrade to VMFS-3**. You must perform this step for each storage device you want to upgrade. ■ **Upgrade the virtual machines** — To upgrade a virtual machine from VMS-2 to VMS-3, right-click the virtual machine in the inventory panel and choose **Upgrade Virtual Machine**.
`esxcfg-vmhbadevs`	Prints a map of VMkernel storage devices to service console devices. There is no VI Client equivalent for this command.
`esxcfg-vmknic`	Creates and updates VMkernel TCP/IP settings for VMotion, NAS, and iSCSI. To set up VMotion, NFS, or iSCSI network connections in VI Client, click **Networking>Add Networking**. Select **VMkernel** and step through the **Add Network Wizard**. Define the IP address subnet mask and VMkernel default gateway in the **Connection Settings** step. To review your settings, click the blue icon to the left of the VMotion, iSCSI, or NFS port. To edit any of these settings, click **Properties** for the switch. Select the port from the list on the switch **Properties** dialog box and click **Edit** to open the port **Properties** dialog box and change the settings for the port. For complete instructions on how to create and update VMotion, NFS, or iSCSI network connections, see "VMkernel Networking Configuration" on page 30.

Table A-1. ESX Server 3 Technical Support Commands (Continued)

Service Console Command	Command Purpose and VI Client Procedure
`esxcfg-vswif`	Creates and updates service console network settings. This command is used if you cannot manage the ESX Server 3 host through the VI Client because of network configuration issues. For more information, see "Troubleshooting Service Console Networking" on page 81. To set up connections for the service console in VI Client, click **Networking>Add Networking**. Select **Service Console** and step through the Add Network Wizard. Define the IP address subnet mask and the service console default gateway in the **Connection Settings** step. To review your settings, click the blue icon to the left of the service console port. To edit any of these settings, click **Properties** for the switch. Select the service console port from the list on the switch **Properties** dialog box. Click **Edit** to open the port **Properties** dialog box and change the settings for the port. For complete instructions on how to create and update the service console connection, see "Service Console Configuration" on page 34.
`esxcfg-vswitch`	Creates and updates virtual machine network settings. To set up connections for a virtual machine in VI Client, click **Networking>Add Networking**. Select **Virtual Machine** and step through the **Add Network Wizard**. To review your settings, click the speech bubble icon to the left of the virtual machine port group. To edit any of these settings, click **Properties** for the switch. Select the virtual machine port from the list on the switch **Properties** dialog box, then click **Edit** to open the port **Properties** dialog box and change the settings for the port. For complete instructions on how to create and update virtual machines, see "Virtual Network Configuration for Virtual Machines" on page 27.

Other Commands

To support certain internal operations, ESX Server 3 installations include a subset of standard Linux configuration commands, for example, network and storage configuration commands. Using these commands to perform configuration tasks can result in serious configuration conflicts and render some ESX Server 3 functions unusable. Always work through the VI Client when configuring ESX Server 3 unless otherwise instructed in VMware Infrastructure documentation or by VMware Technical Support.

Using vmkfstools

B

You use the `vmkfstools` utility to create and manipulate virtual disks, file systems, logical volumes, and physical storage devices on the VMware ESX Server hosts. Using `vmkfstools`, you can create and manage virtual machine file system (VMFS) on a physical partition of a disk. You can also use the command to manipulate files, such as virtual disk files, stored on VMFS-2, VMFS-3, and NFS.

You can perform most `vmkfstools` operations using the VI Client. For information on using the VI Client to work with storage, see "Configuring Storage" on page 103.

This appendix covers the following sections:

- "vmkfstools Command Syntax" on page 283
- "vmkfstools Options" on page 284

vmkfstools Command Syntax

Generally, you don't need to log in as the root user to run the `vmkfstools` commands. However, some commands, such as the file system commands, might require the root user login.

Use the following arguments with the `vmkfstools` command:

- `<options>` are one or more command line options and associated arguments you use to specify the activity for `vmkfstools` to perform — for example, choosing the disk format when creating a new virtual disk.

  After entering the option, specify a file or VMFS file system on which to perform the operation by entering a relative or absolute file path name in the `/vmfs` hierarchy.

- <partition> specifies disk partitions. This argument uses a vmhbaA:T:L:P format, where A, T, L, and P are integers representing adapter, target, LUN, and partition number respectively. The partition digit must be greater than zero (0) and should correspond to a valid VMFS partition of type fb.

  For example, vmhba0:2:3:1 refers to the first partition on LUN 3, target 2, HBA 0.

- <device> specifies devices or logical volumes. This argument uses a path name in the ESX Server 3 device file system. The path name begins with /vmfs/devices, which is the mount point of the device file system.

  Use the following formats when you specify different types of devices:

  - /vmfs/devices/disks for local or SAN-based disks.
  - /vmfs/devices/lvm for ESX Server 3 logical volumes.
  - /vmfs/devices/generic for generic SCSI devices, such as tape drives.

- <path> specifies a VMFS file system or file. This argument is an absolute or relative path that names a directory symbolic link, a raw device mapping, or a file under /vmfs.

  - To specify a VMFS file system, use this format:

    /vmfs/volumes/<file_system_UUID>

    or

    /vmfs/volumes/<file_system_label>

  - To specify a VMFS file, use this format:

    /vmfs/volumes/<file system label|file system UUID>/[dir]/myDisk.vmdk

    You do not need to enter the entire path if the current working directory is the parent directory of myDisk.vmdk.

    For example,

    /vmfs/volumes/datastore1/rh9.vmdk

vmkfstools Options

This section includes a list of all options used with the vmkfstools command. Some of the tasks in this section include options suggested for advanced users only.

The long and short (single letter) forms of options are equivalent. For example, the following commands are identical:

```
vmkfstools --createfs vmfs3 --blocksize 2m vmhba1:3:0:1
vmkfstools -C vmfs3 -b 2m vmhba1:3:0:1
```

-v Suboption

The –v suboption indicates the verbosity level of the command output. The format for this suboption is as follows:

```
–v ––verbose <number>
```

You specify the <number> value as an integer from 1 through 10.

You can specify the –v suboption with any vmkfstools option. If the output of the option isn't suitable for use with the –v suboption, vmkfstools ignores –v.

NOTE Because you can include the –v suboption in any vmkfstools command line, –v is not included as a suboption in the option descriptions.

File System Options

File system options allow you to create a VMFS file system. These options do not apply to NFS. You can perform many of these tasks through the VI Client.

Creating a VMFS File System

```
–C ––createfs vmfs3
        –b ––blocksize <block_size>kK|mM
        –S ––setfsname <fsName>
```

This option creates a VMFS-3 file system on the specified SCSI partition, such as vmhba1:0:0:1. The partition becomes the file system's head partition.

VMFS-2 file systems are read-only on any ESX Server 3 host. You cannot create or modify VMFS-2 file systems but you can read files stored on VMFS-2 file systems. VMFS-3 file systems are not accessible from ESX 2.x hosts.

CAUTION You can have only one VMFS volume for a LUN.

You can specify the following suboptions with the –C option:

- –b ––blocksize – Define the block size for the VMFS-3 file system. The default block size is 1MB. The <block_size> value you specify must be a multiple of 128kb, with a minimum value of 128kb. When you enter a size, indicate the unit type by adding a suffix of m or M. The unit type is not case sensitive — vmkfstools interprets m or M to mean megabytes and k or K to mean kilobytes.

- `-S --setfsname` – Define the volume label of a VMFS volume for the VMFS-3 file system you are creating. Use this suboption only in conjunction with the `-C` option. The label you specify can be up to 128 characters long and cannot contain any leading or trailing blank spaces.

  After you define a volume label, you can use it whenever you specify the VMFS volume for the `vmkfstools` command. The volume label appears in listings generated for the Linux `ls -l` command and as a symbolic link to the VMFS volume under the `/vmfs/volumes` directory.

  To change the VMFS volume label, use the Linux `ln -sf` command. Use the following as an example:

  ```
  ln -sf /vmfs/volumes/<UUID> /vmfs/volumes/<fsName>
  ```

  `<fsName>` is the new volume label to use for the `<UUID>` VMFS.

Example for Creating a VMFS File System

```
vmkfstools -C vmfs3 -b 1m -S my_vmfs /vmfs/devices/disks/vmhba1:3:0:1
```

This example illustrates creating a new VMFS-3 file system named `my_vmfs` on the first partition of target 3, LUN 0 of vmhba adapter 1. The file block size is 1MB.

Extending an Existing VMFS-3 Volume

```
-Z --extendfs <extention-device> <existing-VMFS-volume>
```

This option adds another extent to a previously created VMFS volume `<existing-VMFS-volume>`. You must specify the full path name, for example `/vmfs/devices/disks/vmhba0:1:2:1`, not just the short name `vmhba0:1:2:1`. Each time you use this option, you extend a VMFS-3 volume with a new extent so that the volume spans multiple partitions. At most, a logical VMFS-3 volume can have 32 physical extents.

CAUTION When you run this option, you lose all data that previously existed on the SCSI device you specified in `<extension-device>`.

Example for Extending a VMFS-3 Volume

```
vmkfstools -Z /vmfs/devices/disks/vmhba0:1:2:1
           /vmfs/devices/disks/vmhba1:3:0:1
```

This example extends the logical file system by allowing it to span to a new partition. The extended file system spans two partitions—`vmhba1:3:0:1` and `vmhba0:1:2:1`. In this example, `vmhba1:3:0:1` is the name of the head partition.

Listing Attributes of a VMFS Volume

```
-P --queryfs
        -h --human-readable
```

When you use this option on any file or directory that resides on a VMFS volume, the option lists the attributes of the specified volume. The listed attributes include the VMFS version number (VMFS-2 or VMFS-3), the number of extents comprising the specified VMFS volume, the volume label if any, the UUID, and a listing of the device names where each extent resides.

NOTE If any device backing VMFS file system goes offline, the number of extents and available space change accordingly.

You can specify the `-h` suboption with the `-P` option. If you do so, `vmkfstools` lists the capacity of the volume in a more readable form, for example, `5k`, `12.1M`, or `2.1G`.

Upgrading a VMFS-2 to VMFS-3

You can upgrade a VMFS-2 file system to VMFS-3.

CAUTION The VMFS-2 to VMFS-3 conversion is a one-way process. After you have converted a VMFS-2 volume to VMFS-3, you cannot revert it back to a VMFS-2 volume.

You can upgrade a VMFS-2 file system only if its file block size does not exceed 8 MB.

When upgrading the file system, use the following options:

- `-T --tovmfs3 -x --upgradetype [zeroedthick|eagerzeroedthick|thin]`

  This option converts a VMFS-2 file system VMFS-3 preserving all files on the file system. Before conversion, unload the `vmfs2` and `vmfs3` drivers and load the auxiliary file system driver, `fsaux`, with a module option `fsauxFunction=upgrade`.

  You must specify the upgrade type using the `-x --upgradetype` suboption as one of the following:

  - `-x zeroedthick` (default) – Retains the properties of VMFS-2 thick files. With the `zeroedthick` file format, disk space is allocated to the files for future use and the unused data blocks are not zeroed out.
  - `-x eagerzeroedthick` – Zeroes out unused data blocks in thick files during conversion. If you use this suboption, the upgrade process might take much longer than with the other options.

- `-x thin` – Converts the VMFS-2 thick files into thin-provisioned VMFS-3 files. As opposed to `thick` file format, the thin-provisioned format doesn't allow files to have extra space allocated for their future use, but instead provides the space on demand. During this conversion, unused blocks of the `thick` files are discarded.

  During conversion, the ESX Server 3 file-locking mechanism ensures that no other local process accesses the VMFS volume that is being converted, although you need to make sure that no remote ESX Server host is accessing this volume. The conversion might take several minutes and returns to the command prompt when complete.

  After conversion, unload the `fsaux` driver and load `vmfs3` and `vmfs2` drivers to resume normal operations.

- `-u --upgradefinish`

  This option completes the upgrade.

Virtual Disk Options

Virtual disk options allow you to set up, migrate, and manage virtual disks stored in VMFS-2, VMFS-3, and NFS file systems. You can also perform most of these tasks through the VI Client.

Supported Disk Formats

When you create or clone a virtual disk, you can use the `-d --diskformat` suboption to specify the format for the disk. Choose from the following formats:

- `zeroedthick` (default) – Space required for the virtual disk is allocated during creation. Any data remaining on the physical device is not erased during creation, but is zeroed out on demand at a later time on first write from the virtual machine. The virtual machine does not read stale data from disk.
- `eagerzeroedthick` – Space required for the virtual disk is allocated at creation time. In contrast to `zeroedthick` format, the data remaining on the physical device is zeroed out during creation. It might take much longer to create disks in this format than to create other types of disks.
- `thick` – Space required for the virtual disk is allocated during creation. This type of formatting doesn't zero out any old data that might be present on this allocated space. A non-root user is not allowed to create this format.

- `thin` – Thin-provisioned virtual disk. Unlike with the `thick` format, space required for the virtual disk is not allocated during creation, but is supplied, zeroed out, on demand at a later time.
- `rdm` – Virtual compatibility mode raw disk mapping.
- `rdmp` – Physical compatibility mode (pass-through) raw disk mapping.
- `raw` – Raw device.
- `2gbsparse` – A sparse disk with 2GB maximum extent size. You can use disks in this format with other VMware products, however, you cannot power on sparse disk on an ESX Server host unless you first reimport the disk with `vmkfstools` in a compatible format, such as `thick` or `thin`.
- `monosparse` – A monolithic sparse disk. You can use disks in this format with other VMware products.
- `monoflat` – A monolithic flat disk. You can use disks in this format with other VMware products.

NOTE The only disk formats you can use for NFS are `thin`, `thick`, `zerodthick` and `2gbsparse`.

`Thick`, `zeroedthick` and `thin` usually mean the same because the NFS server and not the ESX Server host decides the allocation policy. The default allocation policy on most NFS servers is `thin`.

Creating a Virtual Disk

```
-c --createvirtualdisk <size>[kK|mM|gG]
        -a --adaptertype [buslogic|lsilogic] <srcfile>
        -d --diskformat [thin|zeroedthick|eagerzeroedthick]
```

This option creates a virtual disk at the specified path on a VMFS volume. Specify the size of the virtual disk. When you enter the value for `<size>`, you can indicate the unit type by adding a suffix of `k` (kilobytes), `m` (megabytes), or `g` (gigabytes). The unit type is not case sensitive—`vmkfstools` interprets either `k` or `K` to mean kilobytes. If you don't specify a unit type, `vmkfstools` defaults to bytes.

You can specify the following suboptions with the `-c` option.

- `-a` specifies the device driver that is used to communicate with the virtual disks. You can choose between BusLogic and LSI Logic SCSI drivers.
- `-d` specifies disk formats. For detailed description of the disk formats, see "Supported Disk Formats" on page 288.

Example for Creating a Virtual Disk

```
vmkfstools –c 2048m /vmfs/volumes/myVMFS/rh6.2.vmdk
```

This example illustrates creating a two-gigabyte virtual disk file named `rh6.2.vmdk` on the VMFS file system named `myVMFS`. This file represents an empty virtual disk that virtual machines can access.

Initializing a Virtual Disk

```
–w ––writezeros
```

This option cleans the virtual disk by writing zeros over all its data. Depending on the size of your virtual disk and the I/O bandwidth to the device hosting the virtual disk, completing this command might take a long time.

CAUTION When you use this command, you lose any existing data on the virtual disk.

Inflating a Thin Virtual Disk

```
–j ––inflatedisk
```

This option converts a `thin` virtual disk to `eagerzeroedthick`, preserving all existing data. The option allocates and zeroes out any blocks that are not already allocated.

See "Supported Disk Formats" on page 288.

Deleting a Virtual Disk

```
–U ––deletevirtualdisk
```

This option deletes files associated with the virtual disk listed at the specified path on the VMFS volume.

Renaming a Virtual Disk

```
–E ––renamevirtualdisk <oldName> <newName>
```

This option renames a file associated with the virtual disk listed in the path specification portion of the command line. You must specify the original file name or file path <`oldName`> and the new file name or file path <`newName`>.

Cloning a Virtual or Raw Disk

```
–i ––importfile <srcfile> –d ––diskformat
    [rdm:<device>|rdmp:<device>|
        raw:<device>|thin|2gbsparse|monosparse|monoflat]
```

This option creates a copy of a virtual disk or raw disk you specify.

You can use the `-d` suboption for the `-i` option. This suboption specifies the disk format for the copy you create. See "Supported Disk Formats" on page 288. A non-root user is not allowed to clone a virtual disk or a raw disk.

NOTE To clone the ESX Server 3 Redo logs while preserving their hierarchy, use the `cp` command.

Example for Cloning a Virtual Disk

```
vmkfstools -i /vmfs/volumes/templates/gold-master.vmdk
              /vmfs/volumes/myVMFS/myOS.vmdk
```

This example illustrates cloning the contents of a master virtual disk from the `templates` repository to a virtual disk file named `myOS.vmdk` on the `myVMFS` file system. You can configure a virtual machine to use this virtual disk by adding lines to the virtual machine configuration file, as in the following example:

```
scsi0:0.present = TRUE
scsi0:0.fileName = /vmfs/volumes/myVMFS/myOS.vmdk
```

Migrating VMware Workstation and VMware GSX Server Virtual Machines

You cannot use a VI Client to migrate virtual machines created with VMware Workstation or VMware GSX Server into your ESX Server 3 system. However, you can use the `vmkfstools -i` command to import the virtual disk into your ESX Server 3 system and then attach this disk to a new virtual machine you create in ESX Server 3. You must import the virtual disk first because you cannot power on disks exported in `2gbsparse` format on an ESX Server host.

To migrate VMware Workstation and GSX Server virtual machines

1. Import a Workstation or GSX Server disk into your `/vmfs/volumes/myVMFS/` directory or any subdirectory.
2. In the VI Client, create a new virtual machine using the **Custom** configuration option.
3. When you configure a disk, select **Use an existing virtual disk** and attach the Workstation or GSX Server disk you imported.

Extending a Virtual Disk

`-X --extendvirtualdisk <newSize>[kK|mM|gG]`

This option extends the size of a disk allocated to a virtual machine after the virtual machine has been created. You must power off the virtual machine that uses this disk file before you enter this command. You might have to update the file system on the disk so the guest operating system can recognize and use the new size of the disk and take advantage of the extra space.

You specify the `newSize` parameter in kilobytes, megabytes, or gigabytes by adding a `k` (kilobytes), `m` (megabytes), or `g` (gigabytes) suffix. The unit type is not case sensitive—`vmkfstools` interprets either `k` or `K` to mean kilobytes. If you don't specify a unit type, `vmkfstools` defaults to kilobytes.

The `newSize` parameter defines the entire new size, not just the increment you add to the disk.

For example, to extend a 4g virtual disk by 1g, enter:

```
vmkfstools -X 5g <disk name>.dsk
```

NOTE Do not extend the base disk of a virtual machine that has snapshots associated with it. If you do, you can no longer commit the snapshot or revert the base disk to its original size.

Migrating a VMFS-2 Virtual Disk to VMFS-3

`-M --migratevirtualdisk`

This option converts the specified virtual disk file from ESX Server 2 format to ESX Server 3 format.

Creating a Virtual Compatibility Mode Raw Device Mapping

`-r --createrdm <device>`

This option creates a Raw Device Mapping (RDM) file on a VMFS-3 volume and maps a raw disk to this file. After this mapping is established, you can access the raw disk as you would a normal VMFS virtual disk. The file length of the mapping is the same as the size of the raw disk it points to.

When specifying the <device> parameter, enter `0` for the partition to indicate that the entire raw disk is used. Use the following format:

`/vmfs/devices/disks/vmhbaA:T:L:0`

See "vmkfstools Command Syntax" on page 283 for more information.

For more details on configuring and using RDMs, see "Raw Device Mapping" on page 145.

NOTE All VMFS-3 file-locking mechanisms apply to RDMs.

Example for Creating a Virtual Compatibility Mode RDM

```
vmkfstools –r /vmfs/devices/disks/vmhba1:3:0:0 my_rdm.vmdk
```

Creates an RDM file named `my_rdm.vmdk` and maps the `vmhba1:3:0:0` raw disk to that file. You can configure a virtual machine to use the `my_rdm.vmdk` mapping file by adding the following lines to the virtual machine configuration file:

```
scsi0:0.present = TRUE
scsi0:0.fileName = /vmfs/volumes/myVMFS/my_rdm.vmdk
```

Creating a Physical Compatibility Mode Raw Device Mapping

```
–z ––createrdmpassthru <device>
```

This option lets you map a pass-through raw device to a file on a VMFS volume. This mapping lets a virtual machine bypass ESX Server 3 SCSI command filtering when accessing its virtual disk.This type of mapping is useful when the virtual machine needs to send proprietary SCSI commands, for example, when SAN-aware software runs on the virtual machine.

After you establish this type of mapping, you can use it to access the raw disk just as you would any other VMFS virtual disk.

When specifying the `<device>` parameter, enter `0` for the partition indicating that the entire raw device is used. Use the following format:

```
/vmfs/devices/disks/vmhbaA:T:L:0
```

See "vmkfstools Command Syntax" on page 283.

Listing Attributes of an RDM

```
–q ––queryrdm
```

This option lets you list the attributes of a raw disk mapping.

This option prints the vmhba name of the raw disk RDM. The option also prints other identification information, like the disk ID, for the raw disk.

Displaying Virtual Disk Geometry

`-g --geometry`

This option gets information about the geometry of a virtual disk.

The output is in the form: `Geometry information C/H/S`, where `C` represents the number of cylinders, `H` represents the number of heads, and `S` represents the number of sectors.

NOTE When you import VMware Workstation virtual disks to an ESX Server 3 host, you might see a disk geometry mismatch error message. A disk geometry mismatch might also be the cause of problems loading a guest operating system or running a newly-created virtual machine.

Managing SCSI Reservations of LUNs

The `-L` option allows you to perform administrative task for physical storage devices. You can perform most of these tasks through the VI Client.

`-L --lock [reserve|release|lunreset|targetreset|busreset]<device>`

This option lets you reserve a SCSI LUN for exclusive use by an ESX Server 3 host, release a reservation so that other hosts can access the LUN, and reset a reservation, forcing all reservations from the target to be released.

CAUTION Using the `-L` option can interrupt the operations of other servers on a SAN. Use the `-L` option only when troubleshooting clustering setups.

Unless specifically advised by VMware, never use this option on a LUN hosting a VMFS volume.

You can specify the `-L` option in several ways:

- `-L reserve` – Reserves the specified LUN. After the reservation, only the server that reserved that LUN can access it. If other servers attempt to access that LUN, a reservation error results.
- `-L release` – Releases the reservation on the specified LUN. Other servers can access the LUN again.
- `-L lunreset` – Resets the specified LUN by clearing any reservation on the LUN and making the LUN available to all servers again. The reset does not affect any of the other LUNs on the device. If another LUN on the device is reserved, it remains reserved.

- `-L targetreset` – Resets the entire target. The reset clears any reservations on all the LUNs associated with that target and makes the LUNs available to all servers again.
- `-L busreset` – Resets all accessible targets on the bus. The reset clears any reservation on all the LUNs accessible through the bus and makes them available to all servers again.

When entering the <device> parameter, use the following format:

`/vmfs/devices/disks/vmhbaA:T:L:P`

See "vmkfstools Command Syntax" on page 283.

Index

Symbols

* next to path **142**

A

accessing storage **95**
adding
  Fibre Channel storage **108**
  groups to ESX Server hosts **224**
  iSCSI hardware-initiated storage **119**
  iSCSI software-initiated storage **125**
  local SCSI storage **104**
  NFS storage **129**
  users to ESX Server hosts **221**
  users to groups **224**
Administrator role **218**
asterisk next to path **142**
authenticating
  groups **215**
  users **214**
authentication daemon **211**

B

Blade servers
  and virtual networking **78**
  configuring a virtual machine port group **78**
  configuring a VMkernel port **79**

C

canonical paths **141**
certificates
  certificate file **225**
  configuring ESX Server searches **228**
  disabling SSL for VI Web Access and SDK **227**
  key file **225**
  location **225**
certification **175**
changing
  password aging for ESX Server **243**
  password aging for users and groups **243**
  proxy services for ESX Server **228**
  service console password plugin **248**
  SSH configuration **253**
checking authentication for iSCSI adapters **206**
CIM and firewall ports **188**
closing ports in the service console firewall **240**
command reference for ESX Server **277**
compatibility modes
  physical **151**
  virtual **151**
configuring
  delegate user **233**
  ESX Server certificate searches **228**
  Fibre Channel storage **108**

hardware-initiated iSCSI storage **119**
local SCSI storage **104**
multipathing for Fibre Channel storage **143**
password complexity **247**
password reuse rules **246**
RDM **155**
software-initiated iSCSI storage **125**
current multipathing state **141**

D

DAS firewall port for ESX Server **183**
datastores
adding extents **136**
and file systems **91**
configuring on NFS volumes **129**
creating on Fibre Channel devices **108**
creating on hardware-initiated iSCSI storage **119**
creating on SCSI disk **104**
creating on software-initiated iSCSI storage **125**
managing **133**
renaming **136**
rescanning **126**
viewing in VI Client **97**
delegate user **232**, **233**
deployments for security **257**
determining the firewall security level for the service console **238**
DHCP **39**
disabling
authentication for iSCSI adapters **208**
cut and paste for guest operating systems **264**
limiting variable information size for guest operating systems **268**
logging for guest operating systems **269**, **273**
SSL for VI Web Access and SDK **227**
DNS **61**
dynamic discovery **112**

E

encryption
and enabling and disabling SSL **225**
for user name, passwords, and packets **225**
ESX Server
adding groups **224**
adding users **221**
architecture and security features **164**
authentication **211**
authentication for iSCSI storage **205**
changing proxy services **228**
cipher strength for connections **249**
command reference **277**
delegate user **232**
deployments and security **257**
host to host firewall ports **188**
password restrictions **241**
security overview **164**
users **211**
virtual switch security **198**
VLAN security **198**
ESX Server host passwords
aging **242**
changing the plugin **248**
complexity **244**
configuring password complexity **247**

configuring password reuse rules **246**
new password criteria **244**
esxcfg commands **277**
EUI identifier **111**
exporting ESX Server host users and groups **220**
extents **136**

F

failover **55**
failover paths
status **142**
Fibre Channel storage
adding **108**
overview **107**
file systems
managing **133**
NFS **91**
upgrading **135**
VMFS **91**
firewall ports
and encryption **225**
backup agents **237**
CIM **188**
configured with a VirtualCenter Server **179**
configured without a VirtualCenter Server **182**
determining the firewall security level for the service console **238**
for connecting the virtual machine console **186**
for management access **183**
FTP **188**
host to host **188**
iSCSI software client **188**
license server and VirtualCenter Server **179**
management **188**
NFS **188**
NIS **188**
opening and closing for the service console **240**
opening with the VI Client **188**
overview **177**
SDK and the virtual machine console **186**
security level **237**
service console **237**
setting the security level for the service console firewall **239**
SMB **188**
SNMP **188**
SSH **188**
supported services **188**
VI Client and the virtual machine console **186**
VI Client and VirtualCenter Server **179**
VI Client direct connection **182**
VI Web Access and the virtual machine console **186**
VI Web Access and VirtualCenter Server **179**
VI Web Access direct connection **182**
Fixed path policy **137**
FTP and firewall ports **188**

G

groups
adding to ESX Server hosts **224**
authentication **215**
exporting a group list **220**
Groups table for ESX Server hosts **219**
modifying on ESX Server hosts **224**

removing from ESX Server hosts **225**
viewing group lists **220**
guest operating systems
disabling cut and paste **264**
disabling logging **269**, **273**
limiting variable information size **268**
security recommendations **263**

H

HTTP and HTTPS firewall port **183**

I

IQN identifier **111**
iSCSI
authenticating **205**
CHAP **205**
checking authentication **206**
configuring CHAP authentication **207**
disabling authentication **208**
firewall port for ESX Server **183**
networking **73**
protecting transmitted data **208**
QLogic iSCSI adapters **204**
security **204**
software client and firewall ports **188**
iSCSI hardware-initiated storage
adding **119**
overview **113**
iSCSI HBA
alias **115**
CHAP authentication **118**
CHAP parameters **115**
dynamic discovery **115**
static discovery **115**
iSCSI networking
creating a service console connection **76**
creating a VMkernel port **74**
iSCSI securing ports **208**
iSCSI software-initiated storage
adding **125**
overview **121**
iSCSI storage
discovery methods **112**
EUI identifier **111**
hardware-initiated **110**
initiators **110**
IQN identifier **111**
name formats **111**
security **112**
software-initiated **110**
isolation
virtual machine **165**
virtual networking layer **169**
virtual switches **169**
VLANs **169**

L

Layer 2 security **51**
license server
firewall ports for **183**
firewall ports with VirtualCenter Server **179**
load balancing **55**
local SCSI storage
adding **104**
overview **104**

M

MAC address
configuring **68**
generating **67**
Manage Paths wizard **144**

management access
firewall ports **183**
modifying
groups on ESX Server hosts **224**
users on ESX Server hosts **223**
Most Recently Used path policy **137**
multipathing
active paths **142**
canonical paths **141**
dead paths **142**
disabled paths **142**
failover **143**
managing **143**
standby paths **142**
multipathing policy
setting **143**
multipathing state **141**

N

NAS
firewall port for ESX Server **183**
mounting **70**
Nessus **254**
networking best practices **69**
networks
security **195**
NFS
delegate users **232**
firewall ports **188**
NFS storage
adding **129**
overview **127**
NIC teaming
definition **20**
NIS and firewall ports **188**
No Access role **218**

O

opening ports in the service console firewall **240**

P

pam_cracklib.so plugin **244**
pam_passwdqc.so plugin **248**
password restrictions
aging **242**
complexity **244**
for the ESX Server host **241**
minimum length **244**
path failure **137**
path policies
Fixed **137**
Most Recently Used **137**
Round Robin **138**
paths
preferred **142**, **144**
permissions
and privileges **215**
overview **215**
root user **215**
VirtualCenter administrator **215**
vpxuser **215**
port group
configuring **58**
definition **20**
using **24**
preferred path **142**, **144**
preventing malicious device disconnection **266**
privileges
and permissions **215**
proxy services
and encryption **225**
changing **228**

R

raw device mapping
see RDM **145**
RDM
advantages **147**
and virtual disk files **154**
and vmkfstools **158**
creating **155**
dynamic name resolution **152**
overview **145**
physical compatibility mode **151**
virtual compatibility mode **151**
with clustering **154**
Read Only role **218**
removing
groups from ESX Server hosts **225**
users from ESX Server hosts **224**
users from groups **224**
resource guarantees and security **165**
resource limits and security **165**
roles
Administrator **218**
and permissions **218**
default **218**
No Access **218**
Read Only **218**
root login
delegate user **232**
permissions **215**
SSH **253**
Round Robin path policy **138**
routing **61**
RPMs **254**

S

SCSI
vmkfstools **283**
SDK and firewall ports for connecting to the virtual machine console **186**
security
CHAP authentication **205**
cipher strength **249**
delegate user **232**
direct access users **214**
encryption **225**
ESX Server architecture **164**
example, DMZ in a single ESX Server host **170**, **172**
forged transmissions **202**
groups **215**
iSCSI storage **204**
MAC address changes **202**
overview of users, groups, permissions, and roles **213**
PAM authentication **211**
password restrictions for the ESX Server host **241**
patches **254**
permissions **215**
promiscuous mode **202**
recommendations for virtual machines **263**
roles **218**
scanning software **254**
security certificates **225**
service console firewall security level **237**
service console security measures **167**
setgid applications **250**
setuid applications **250**
SSH connections **253**
user authentication **211**

user management **211**
virtual machines **165**
virtual network **195**
virtual networking layer **169**
VirtualCenter users **214**
virtualization layer **164**
VLAN hopping **198**
VLANs **195**
VMkernel **164**
VMware policy **175**
vmware-authd **211**
service console
direct connections **237**
logging on **237**
password restrictions **241**
recommendations for securing **236**
remote connections **237**
securing with VLANs and virtual switches **198**
security **167**
setgid applications **250**
setuid applications **250**
SSH connections **253**
service console networking
configuration **34**
troubleshooting **81**
setgid applications **250**
setting security level for the service console firewall **239**
setting up CHAP authentication for iSCSI adapters **207**
setuid applications **250**
single point of failure **104**
SMB and firewall ports **188**
SNMP and firewall ports **188**
SPOF **104**
SSH
changing configuration **253**
firewall ports **188**
security settings **253**
static discovery **112**
storage
access for virtual machines **95**
adapters **91**
configuration tasks **101**
Fibre Channel **107**
iSCSI **110**
local SCSI **104**
NFS **127**
SAN **107**
securing with VLANs and virtual switches **198**
types **88**
viewing in VI Client **97**
storage adapters
Fibre Channel **107**
iSCSI HBA **115**
rescanning **126**
viewing in VI Client **99**

T

TCP ports **183**
third-party software support policy **175**
Tomcat Web service **167**
traffic shaping **53**

U

UDP ports **183**
users
adding to ESX Server hosts **221**
authentication **214**
direct access users **214**
exporting a user list **220**
from Windows domain **214**
modifying on ESX Server hosts **223**

removing from ESX Server hosts **224**
Users table for ESX Server hosts **219**
viewing user list **220**
VirtualCenter users **214**

V

VI Client
firewall ports for connecting to the virtual machine console **186**
firewall ports for direct connection **182**
firewall ports with VirtualCenter Server **179**
VI Web Access
and ESX Server services **225**
disabling SSL **227**
firewall ports for connecting to the virtual machine console **186**
firewall ports for direct connection **182**
firewall ports with VirtualCenter Server **179**
viewing ESX Server host users and groups **220**
virtual machine networking **27**
virtual machines
configuring a delegate user **233**
delegate user **232**
disabling copy and paste **264**
disabling logging **269**, **273**
isolation example **170**, **172**
limiting variable information size **268**
preventing device disconnection **266**
resource reservations and limits **165**
security **165**
security recommendations **263**
virtual networking layer and security **169**
virtual switches
802.1Q and ISL tagging attacks **200**
and iSCSI **208**
double-encapsulated attacks **200**
forged transmissions **202**
MAC address changes **202**
MAC flooding **200**
multicast brute-force attacks **200**
promiscuous mode **202**
random frame attacks **200**
scenarios for deployment **257**
security **200**
spanning tree attacks **200**
VirtualCenter Server
firewall ports **179**
permissions **215**
virtualization layer and security **164**
VLAN
definition **20**
VLANs
and iSCSI **208**
Layer 2 security **198**
scenarios for deployment **257**
security **195**
VLAN hopping **198**
VMFS
sharing **257**
vmkfstools **283**
VMkernel
configuring **30**
definition **20**
security **164**

vmkfstools
file system options **285**
overview **283**
syntax **283**
virtual disk options **288**
VMotion
definition **20**
firewall port **183**
networking configuration **30**
securing with VLANs and virtual switches **198**
vSwitch
definition **20**
policies **50**
using **21**

Updates for the ESX Server 3 Configuration Guide

Last Updated: May 14, 2010

This document provides updates to the Update 2 release of the *ESX Server 3 Configuration Guide* for ESX Server 3.5 and VirtualCenter 2.5. Updated descriptions, procedures, and graphics are organized by page number so that you can easily locate the areas of the guide that have changes. If the change spans multiple sequential pages, this document provides the starting page number only.

The following are the updates to the *ESX Server 3 Configuration Guide*:

- Updates for the Table of TCP and UPD Ports for Management Access on Page 185
- Update for Description of SSL Behavior on Page 225
- Update for Types of Certificates to Install for Certificate Checking on Page 226
- Update for Using Third-Party Applications to Upload Certificates on Page 227
- Update for Certificates Supported for Encryption Over an SSL Connection on Page 227
- Updates to the Password Aging Section on Page 242
- Updates for the To change default password complexity for the pam_cracklib.so plug-in Procedure on Page 247

Updates for the Table of TCP and UPD Ports for Management Access on Page 185

The information in Table 10-1 for ports 5988 and 5989 is incorrect. Port 5988 is no longer required for CIM traffic and should be closed, so information related to Port 5988 should be deleted from the table.

Port 5989 is open for CIM traffic on HTTPS instead of HTTP. The following row replaces the information given for port 5989.

5989	CIM XML transactions over HTTPS.	Incoming and outgoing TCP

Update for Description of SSL Behavior on Page 225

The Encryption and Security Certificates for ESX Server 3 section incorrectly describes SSL behavior for ESX Server 3 and states that SSL is not enabled by default. The corrected paragraph should appear as follows:

SSL certificates are used to encrypt network traffic, but the certificate used for encryption is not verified by default. Therefore, the connection between the VI Client and VirtualCenter is vulnerable to a possible man-in-the-middle attack. To prevent such an attack and to fully enable the security provided by certificates in ESX Server 3, you must enable certificate checking and install new certificates.

Update for Types of Certificates to Install for Certificate Checking on Page 226

The paragraph after the To enable certificate checking procedure, describes how to receive the full benefits of certificate checking, but does not specify what type of certificates to install. The following information should be included:

To receive the full benefit of certificate checking, install new certificates that are signed by a valid internal or public CA.

Update for Using Third-Party Applications to Upload Certificates on Page 227

The Adding Certificates and Modifying ESX Server 3 Web Proxy Settings section states incorrectly that you should avoid setting up certificates by using pass phrases. The section should state the following:

Do not set up certificates by using pass phrases. ESX Server 3 does not support pass phrases, also known as encrypted keys. If you set up a pass phrase, ESX Server 3 processes cannot start correctly.

Update for Certificates Supported for Encryption Over an SSL Connection on Page 227

The Adding Certificates and Modifying ESX Server 3 Web Proxy Settings section omits a note about what type of certificate is supported by ESX Server for encrypting session information sent over an SSL connection. The section should include the following information:

ESX Server supports only X.509 certificates to encrypt session information sent over SSL connections between server and client components.

Updates to the Password Aging Section on Page 242

The Password Aging section incorrectly states that the maximum number of days that a user can retain a password is 90 days. The corrected list item should appear as follows:

Maximum days – The number of days that a user can retain a password. By default, passwords are set to never expire.

Updates for the To change default password complexity for the pam_cracklib.so plug-in Procedure on Page 247

In the To change default password complexity for the pam_cracklib.so plug-in procedure, the information about credits and the example for `esxcfg-auth --usecrack` command are incorrect. The content should be as follows:

- `<lc_credit>` represents lowercase letters
- `<uc_credit>` represents uppercase letters
- `<d_credit>` represents digits
- `<oc_credit>` represents special characters, such as underscore or dash

Credits add to a password's complexity score. A user's password must meet or exceed the minimum score, which you define using the `<minimum_length>` parameter.

NOTE The `pam_cracklib.so` plug-in does not accept passwords that are less than six characters, regardless of credits used and regardless of the value you assign to `<minimum_length>`. In other words, if `<minimum_length>` is 5, users must still enter no fewer than six characters.

To determine whether or not a password is acceptable, the `pam_cracklib.so` plug-in uses several rules to calculate the password score.

- Each character in the password, regardless of the type, counts as one against `<minimum_length>`.
- Nonzero values in the credit parameters affect password complexity differently depending on whether negative or positive values are used.
  - For positive values, add one credit for the character class, up to the maximum number of credits specified by the credit parameter.

    For example, if `<lc_credit>` is 1, add one credit for using a lowercase letter in the password. In this case, one is the maximum number of credits allowed for lowercase letters, regardless of how many are used.

- For negative values, do not add credit for the character class, but mandate that the character class is used a minimum number of times. The minimum number is specified by the credit parameter.

  For example, if <uc_credit> is -1, passwords must contain at least one uppercase character. In this case, no extra credit is given for using uppercase letters, regardless of how many are used.

- Character classes with a value of zero count toward the total length of the password, but do not receive extra credit, nor are they required. You can set all character classes to zero to enforce password length without considering complexity.

  For example, the passwords xyzpqets and Xyzpq3#s would each have a password score of eight.

The plug-in then compares the total score, or effective length, of the password to the value of <minimum_length>.

Example of esxcfg-auth --usecrack Command

```
esxcfg-auth --usecrack=3 9 1 -1 -1 1
```

- Users are allowed three attempts to enter their password before they are locked out.
- The password score must be 9.
- Up to one credit is given for using lowercase letters.
- At least one uppercase letter is required. No extra credit is given for this character class.
- At least one digit is required. No extra credit is given for this character class.
- Up to one credit is given for using special characters.

Using these sample values, the password candidate **xyzpqe#** would fail. While the password score is 9, (x + y + z + p +q + e + #) + (lc_credit + oc_credit) = 9, it does not contain the required uppercase letter and digit.

The password candidate Xyzpq3# would be accepted. The password score for this example is also 9, (X + y + z + p +q + 3 + #) + (lc_credit + oc_credit) = 9, but this password includes the required uppercase letter and digit. The uppercase letter and digit do not add extra credit.