

ESX Configuration Guide

ESX 4.0
vCenter Server 4.0

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-000106-09

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

Copyright © 2009–2011 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at <http://www.vmware.com/go/patents>.

VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

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

Contents

Updated Information 7

About This Book 9

1 Introduction to ESX Configuration 11

Networking

2 Introduction to Networking 15

Networking Concepts Overview 15

Network Services 16

View Networking Information in the vSphere Client 16

View Network Adapter Information in the vSphere Client 17

3 Basic Networking with vNetwork Standard Switches 19

vNetwork Standard Switches 19

Port Groups 20

Port Group Configuration for Virtual Machines 20

VMkernel Networking Configuration 21

Service Console Configuration 23

vNetwork Standard Switch Properties 25

4 Basic Networking with vNetwork Distributed Switches 29

vNetwork Distributed Switch Architecture 29

Configuring a vNetwork Distributed Switch 30

dvPort Groups 32

Private VLANs 34

Configuring vNetwork Distributed Switch Network Adapters 35

Configuring Virtual Machine Networking on a vNetwork Distributed Switch 39

5 Advanced Networking 41

Internet Protocol Version 6 41

Networking Policies 42

Change the DNS and Routing Configuration 55

MAC Addresses 55

TCP Segmentation Offload and Jumbo Frames 57

NetQueue and Networking Performance 60

VMDirectPath Gen I 60

6 Networking Best Practices, Scenarios, and Troubleshooting 63

- Networking Best Practices 63
- Mounting NFS Volumes 64
- Networking Configuration for Software iSCSI Storage 64
- Configuring Networking on Blade Servers 65
- Troubleshooting 67

Storage**7 Introduction to Storage 71**

- About ESX Storage 71
- Types of Physical Storage 72
- Supported Storage Adapters 73
- Target and Device Representations 73
- About ESX Datastores 76
- Comparing Types of Storage 78
- Viewing Storage Information in the vSphere Client 79

8 Configuring ESX Storage 83

- Local SCSI Storage 83
- Fibre Channel Storage 84
- iSCSI Storage 84
- Storage Refresh and Rescan Operations 94
- Create VMFS Datastores 95
- Network Attached Storage 95
- Creating a Diagnostic Partition 97

9 Managing Storage 99

- Managing Datastores 99
- Changing VMFS Datastore Properties 101
- Managing Duplicate VMFS Datastores 103
- Using Multipathing with ESX 105
- Thin Provisioning 113
- TurnOff vCenter Server Storage Filters 116

10 Raw Device Mapping 119

- About Raw Device Mapping 119
- Raw Device Mapping Characteristics 122
- Managing Mapped LUNs 126

Security**11 Security for ESX Systems 131**

- ESX Architecture and Security Features 131
- Security Resources and Information 139

12	Securing an ESX Configuration	141
	Securing the Network with Firewalls	141
	Securing Virtual Machines with VLANs	150
	Securing Virtual Switch Ports	155
	Securing iSCSI Storage	157
13	Authentication and User Management	161
	Securing ESX Through Authentication and Permissions	161
	Encryption and Security Certificates for ESX	168
14	Service Console Security	175
	General Security Recommendations	175
	Log In to the Service Console	176
	Service Console Firewall Configuration	176
	Password Restrictions	179
	Cipher Strength	185
	setuid and setgid Flags	185
	SSH Security	187
	Security Patches and Security Vulnerability Scanning Software	188
15	Security Deployments and Recommendations	191
	Security Approaches for Common ESX Deployments	191
	Virtual Machine Recommendations	195

Host Profiles

16	Managing Host Profiles	203
	Host Profiles Usage Model	203
	Access Host Profiles View	204
	Creating a Host Profile	204
	Export a Host Profile	205
	Import a Host Profile	205
	Edit a Host Profile	206
	Manage Profiles	207
	Checking Compliance	210

Appendixes

A	ESX Technical Support Commands	215
B	Linux Commands Used with ESX	219
C	Using vmkfstools	221
	vmkfstools Command Syntax	221
	vmkfstools Options	222

	Index	231
--	--------------	------------

Updated Information

This *ESX Configuration Guide* is updated with each release of the product or when necessary.

This table provides the update history of the *ESX Configuration Guide*.

Revision	Description
EN-000106-09	Modified a word in “Handling Datastore Over-Subscription,” on page 115.
EN-000106-08	■ Updated procedure in “Generate New Certificates for the ESX Host,” on page 169. ■ Updated procedure in “Replace a Default Certificate with a CA-Signed Certificate,” on page 169.
EN-000106-07	In “Comparing Types of Storage,” on page 78, removed VM Cluster from supported vSphere features, and included citation for Microsoft clustering.
EN-000106-06	Removed a note from the topic “Networking Configuration for Software iSCSI Storage,” on page 64 that indicated software initiated iSCSI was not supported. This feature has been supported since the ESX 4.0 initial product release.
EN-000106-05	■ “Add a Virtual Machine Port Group,” on page 20 reflects that the network label for a port group can contain up to 63 characters. ■ The Sample vSphere Network Configuration and Traffic Flow figure in “Firewalls for Configurations with vCenter Server,” on page 142 now reflects the correct functionality of port 5989.
EN-000106-04	The topic “vNetwork Standard Switches,” on page 19 now refers to <i>Configuration Maximums for vSphere 4.0</i> for port and port group maximums.
EN-000106-03	Appendix A, “ESX Technical Support Commands,” on page 215 now correctly lists the <code>esxcfg-scsidevs</code> command.
EN-000106-02	The topic “Change Default Password Complexity for the pam_cracklib.so Plug-In,” on page 183 now correctly describes the behavior of the <code>pam_cracklib.so</code> plug-in.
EN-000106-01	The topic “Configure a Password Reuse Rule,” on page 182 now correctly lists the file where you can set the number of old passwords that are stored for a user. The correct file is <code>system-auth-generic</code> .
EN-000106-00	Initial release.

About This Book

This manual, the *ESXESXiConfiguration Guide*, provides information on how to configure networking for VMware® ESXESXi, including how to create virtual switches and ports and how to set up networking for virtual machines, VMware VMotion™, IP storage, and the service console. It also discusses configuring the file system and various types of storage such as iSCSI and Fibre Channel. The guide provides a discussion of security features built into ESXESXi and the measures that you can take to safeguard ESXESXi from attack. In addition, it includes a list of ESXESXi technical support commands along with their VMware vSphere™ Client equivalents and a description of the vmkfstools utility.

This information covers ESX 4.0.

Intended Audience

This manual is intended for anyone who needs to install, upgrade, or use ESX. The information in this manual 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 vSphere Documentation

The vSphere documentation consists of the combined VMware vCenter Server and ESX documentation set.

Abbreviations Used in Figures

The figures in this manual use the abbreviations listed in [Table 1](#).

Table 1. Abbreviations

Abbreviation	Description
database	vCenter Server database
datastore	Storage for the managed host
dsk#	Storage disk for the managed host
hostn	vCenter Server managed hosts
SAN	Storage area network type datastore shared between managed hosts
tplt	Template

Table 1. Abbreviations (Continued)

Abbreviation	Description
user#	User with access permissions
VC	vCenter Server
VM#	Virtual machines on a managed host

Technical Support and Education Resources

The following technical support resources are available to you. To access the current version of this book and other books, go to <http://www.vmware.com/support/pubs>.

Online and Telephone Support

To 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

To 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>.

Introduction to ESX Configuration

This guide describes the tasks you need to complete to configure ESX host networking, storage, and security. In addition, it provides overviews, recommendations, and conceptual discussions to help you understand these tasks and how to deploy a host to meet your needs.

Before you use this information, read the *Introduction to vSphere* for an overview of system architecture and the physical and virtual devices that make up a vSphere system.

This introduction summarizes the contents of this guide.

Networking

The networking information provides 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 host's network connections, and a discussion of advanced networking topics and tasks.

Storage

The storage information provides you with a basic understanding of storage, a description of the basic tasks you perform to configure and manage your ESX host's storage, and a discussion of how to set up raw device mapping (RDM).

Security

The security information discusses safeguards that VMware has built into ESX and measures that you can take to protect your 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.

Host Profiles

This section describes the host profiles feature and how it is used to encapsulate the configuration of a host into a host profile. This section also describes how to apply this host profile to another host or cluster, edit a profile, and check a host's compliance with a profile.

Appendixes

The appendixes provide specialized information you might find useful when configuring an ESX host.

- **ESX Technical Support Commands** – Discusses the ESX 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 configuration commands. This appendix provides you with vSphere 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

Introduction to Networking

This introduction to networking guides you through the basic concepts of ESX networking and how to set up and configure a network in a vSphere environment.

This chapter includes the following topics:

- [“Networking Concepts Overview,”](#) on page 15
- [“Network Services,”](#) on page 16
- [“View Networking Information in the vSphere Client,”](#) on page 16
- [“View Network Adapter Information in the vSphere Client,”](#) on page 17

Networking Concepts Overview

A few concepts are essential for a thorough understanding of virtual networking. If you are new to ESX, it is helpful to review these concepts.

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 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 when you add a network.

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 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.

A vNetwork Distributed Switch acts as a single vSwitch across all associated hosts on a datacenter. This allows virtual machines to maintain consistent network configuration as they migrate across multiple hosts.

A dvPort is a port on a vNetwork Distributed Switch that connects to a host’s service console or VMkernel or to a virtual machine’s network adapter.

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. Typically, a single vSwitch is associated with one or more port groups.

A dvPort group is a port group associated with a vNetwork Distributed Switch and specifies port configuration options for each member port. dvPort Groups define how a connection is made through the vNetwork Distributed Switch to the network.

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.

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.

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.

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

Network Services

A vNetwork provides several different services to the host and virtual machines.

You can enable three types of network services in ESX:

- Connecting virtual machines to the physical network and to each other.
- Connecting VMkernel services (such as NFS, iSCSI, or VMotion) to the physical network.
- Running management services for ESX via the service console. A service console port, which is set up by default during installation, is required for ESX to connect to any network or remote services, including the vSphere Client. Additional service console ports might be necessary for other services, such as iSCSI storage.

View Networking Information in the vSphere Client

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

Procedure

- 1 Log in to the vSphere Client and select the host from the inventory panel.
- 2 Click the **Configuration** tab and click **Networking**.
- 3 Click **Virtual Switch** to view vNetwork Standard Switch networking on the host or **Distributed Virtual Switch** to view vNetwork Distributed Switch networking on the host.

The **Distributed Virtual Switch** option appears only on hosts that are associated with a vNetwork Distributed Switch.

Networking information is displayed for each virtual switch on the host.

View Network Adapter Information in the vSphere Client

For each physical network adapter on the host, you can view information such as the speed, duplex, and observed IP ranges.

Procedure

- 1 Log in to the vSphere Client and select the host from the inventory panel.
- 2 Click the **Configuration** tab, and click **Network Adapters**.

The network adapters panel displays the following information.

Table 2-1. Network Panel Information

Option	Description
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 the LAN

Basic Networking with vNetwork Standard Switches

3

The following topics guide you through basic vNetwork Standard Switch (vSwitch) network setup and configuration in a vSphere environment.

- Virtual machines
- VMkernel
- Service console

This chapter includes the following topics:

- [“vNetwork Standard Switches,”](#) on page 19
- [“Port Groups,”](#) on page 20
- [“Port Group Configuration for Virtual Machines,”](#) on page 20
- [“VMkernel Networking Configuration,”](#) on page 21
- [“Service Console Configuration,”](#) on page 23
- [“vNetwork Standard Switch Properties,”](#) on page 25

vNetwork Standard Switches

You can create abstracted network devices called vNetwork Standard Switches (vSwitches). A vSwitch can route traffic internally between virtual machines and link to external networks.

You can use vSwitches to combine the bandwidth of multiple network adapters and balance communications traffic among them. You can also configure a vSwitch to handle physical NIC failover.

A vSwitch models a physical Ethernet switch. The default number of logical ports for a vSwitch is 56. 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. For information on maximum allowed ports and port groups, see *Configuration Maximums for vSphere 4.0*.

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.

Port Groups

Port groups aggregate multiple ports under a common configuration and provide a stable anchor point for virtual machines connecting to labeled networks. You can create a maximum of 512 port groups on a single host.

Each port group is identified by a network label, which is unique to the current host. Network labels are used to make virtual machine configuration portable across hosts. All port groups in a datacenter that are physically connected to the same network (in the sense that each can receive broadcasts from the others) are given the same label. Conversely, if two port groups cannot receive broadcasts from each other, they have distinct labels.

A VLAN ID, which restricts port group traffic to a logical Ethernet segment within the physical network, is optional. For a port group to reach port groups located on other VLANs, the VLAN ID must be set to 4095. If you use VLAN IDs, you must change the port group labels and VLAN IDs together so that the labels properly represent connectivity.

Port Group Configuration for Virtual Machines

You can add or modify a virtual machine port group from the vSphere Client.

The vSphere Client Add Network wizard guides you through the tasks to create a virtual network to which virtual machines can connect, including creating a vSwitch and configuring settings for a network label.

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

ESX does not 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 in the new 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.

Add a Virtual Machine Port Group

Virtual machine port groups provide networking for virtual machines.

Procedure

- 1 Log in to the vSphere Client and select the host from the inventory panel.
- 2 Click the **Configuration** tab and click **Networking**.
- 3 Select the Virtual Switch view.
vSwitches appear in an overview that includes a details layout.
- 4 On the right side of the page, click **Add Networking**.
- 5 Accept the default connection type, **Virtual Machines**, and click **Next**.
- 6 Select **Create a virtual switch** or one of the listed existing vSwitches and the associated physical adapters to use for this port group.

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.

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

NOTE The network label can contain up to 63 characters.

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

- 9 (Optional) If you are using a VLAN, for **VLAN ID**, enter a number between 1 and 4094. If you are not using a VLAN, leave this blank.

If you enter 0 or leave the option 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**.

VMkernel Networking Configuration

A VMkernel networking interface is used for VMware VMotion and IP storage.

Moving a virtual machine from one host to another is called migration. Using VMotion, you can migrate powered on 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. Because these storage types are network based, they can use the same VMkernel interface and port group.

The network services that the VMkernel provides (iSCSI, NFS, and VMotion) 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 provides networking support in multiple ways for each of the services it handles.

The VMkernel TCP/IP stack handles 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

If you have two or more physical NICs for iSCSI, you can create multiple paths for the software iSCSI by using the port binding technique. For more information on port binding, see the *iSCSI SAN Configuration Guide*.

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

Set Up VMkernel Networking

Create a VMkernel network adapter for use as a VMotion interface or an IP storage port group.

Procedure

- 1 Log in to the vSphere Client and select the host from the inventory panel.
- 2 Click the **Configuration** tab and click **Networking**.
- 3 In the Virtual Switch view, click **Add Networking**.
- 4 Select **VMkernel** and click **Next**.
- 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.

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.

- 7 Click **Next**.
- 8 Select or enter a network label and a VLAN ID.

Option	Description
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 host 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 host. If this property is not enabled for any port group, migration with vMotion to this host is not possible.

- 10 Choose whether to use this port group for fault tolerance logging, and click **Next**.
- 11 Select **Obtain IP settings automatically** to use DHCP to obtain IP settings, or select **Use the following IP settings** to specify IP settings manually.

If you choose to specify IP settings manually, provide this information.

- a Enter the IP address and subnet mask for the VMkernel interface.
This address must be different from the IP address set for the service console.
- b Click **Edit** to set the VMkernel Default Gateway for VMkernel services, such as vMotion, NAS, and iSCSI.
- c On the **DNS Configuration** tab, the name of the host is entered by default.
The DNS server addresses that were specified during installation are also preselected, as is the domain.
- d On the **Routing** tab, the service console and the VMkernel each need their own gateway information.
A gateway is needed for connectivity to machines not on the same IP subnet as the service console or VMkernel. The default is static IP settings.
- e Click **OK**, then click **Next**.

- 12 On an IPv6-enabled host, select **No IPv6 settings** to use only IPv4 settings on the VMkernel interface, or select **Use the following IPv6 settings** to configure IPv6 for the VMkernel interface.
This screen does not appear when IPv6 is disabled on the host.
- 13 If you choose to use IPv6 for the VMkernel interface, select one of the following options for obtaining IPv6 addresses.
 - **Obtain IPv6 addresses automatically through DHCP**
 - **Obtain IPv6 addresses automatically through router advertisement**
 - **Static IPv6 addresses**
- 14 If you choose to use static IPv6 addresses, complete the following steps.
 - a Click **Add** to add a new IPv6 address.
 - b Enter the IPv6 address and subnet prefix length, and click **OK**.
 - c To change the VMkernel default gateway, click **Edit**.
- 15 Click **Next**.
- 16 Review the information, click **Back** to change any entries, 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.

Common service console configuration modifications include changing NICs and changing the settings for a NIC that is in use.

If there is only one service console connection, 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.

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

Set Up Service Console Networking

A single service console network interface is set up during the ESX installation process. You can also add additional service console interfaces after ESX is up and running.

Procedure

- 1 Log in to the vSphere Client and select the host from the inventory panel.
- 2 Click the **Configuration** tab, and click **Networking**.
- 3 In the Virtual Switch view, click **Add Networking**.
- 4 Select **Service Console**, 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 network adapters in the system are being used by existing vSwitches.
- 6 Enter the network label and VLAN ID, and click **Next**.
- 7 Enter the IP address and subnet mask, or select **Obtain IP setting automatically**.
- 8 Click **Edit** to set the service console default gateway and click **Next**.

- 9 On an IPv6-enabled host, select **No IPv6 settings** to use only IPv4 settings for the service console, or select **Use the following IPv6 settings** to configure IPv6 for the service console.

This screen does not appear if IPv6 is disabled on the host.

- 10 If you choose to use IPv6, select how to obtain IPv6 addresses.
- 11 If you chose **Static IPv6 addresses**, do the following:
 - a Click **Add** to add a new IPv6 address.
 - b Enter the IPv6 address and subnet prefix length, and click **OK**.
 - c To change the service console default gateway, click **Edit**.
- 12 Click **Next**.
- 13 Review the information, click **Back** to change any entries, and click **Finish**.

Configure Service Console Ports

You can edit service console port properties, such as IP settings and networking policies.

Procedure

- 1 Log in to the vSphere Client and select the host from the inventory panel.
- 2 Click the **Configuration** tab, and click **Networking**.
- 3 On the right side of the page, click **Properties** for the vSwitch that you want to edit.
- 4 In the vSwitch Properties dialog box, click the **Ports** tab.
- 5 Select **Service Console** and click **Edit**.
- 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**.

Set the Default Gateway

You can configure one default gateway for the service console per TCP/IP stack.



CAUTION Make sure that your network settings are correct before saving your changes. If the network settings are misconfigured, the UI can lose connectivity to the host, and you must then reconfigure the host from the command line at the service console.

Procedure

- 1 Log into the vSphere Client and select the host from the inventory panel.
- 2 Click the **Configuration** tab, and click **DNS and Routing**.
- 3 Click **Properties**.
- 4 Click the **Routing** tab.
- 5 Under Service Console, set the default gateway and gateway device for service console networking.

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 is used for the default route.

The service console and VMkernel are often not connected to the same network, so 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.

On an IPv6-enabled host, you can also select a default gateway for IPv6 and a gateway device for IPv6 for service console networking.

- 6 Under VMkernel, set the default gateway for VMkernel networking.

On an IPv6-enabled host, you can also select a default gateway for IPv6 for VMkernel networking.

- 7 Click **OK**.

Display Service Console Information

You can view service console network information, such as the VLAN ID and network policies.

Procedure

- 1 Click the info icon to the left of the service console port group 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, you 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 name to its DHCP-generated IP address, use the service console's numeric IP address to access 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.

vNetwork Standard Switch Properties

vNetwork Standard Switch settings control vSwitch-wide defaults for ports, which can be overridden by port group settings for each vSwitch. You can edit vSwitch properties, such as the uplink configuration and the number of available ports.

Change the Number of Ports for a vSwitch

A vSwitch serves as a container for port configurations that use a common set of network adapters, including sets that contain no network adapters at all. Each virtual switch provides a finite number of ports through which virtual machines and network services can reach one or more networks.

Procedure

- 1 Log in to the vSphere Client and select the host from the inventory panel.
- 2 Click the **Configuration** tab and click **Networking**.
- 3 On the right side of the page, click **Properties** for the vSwitch that you want to edit.
- 4 Click the **Ports** tab.
- 5 Select the vSwitch item in the Configuration list, and click **Edit**.
- 6 Click the **General** tab.
- 7 Choose the number of ports that you want to use from the drop-down menu.
- 8 Click **OK**.

What to do next

Changes will not take effect until the system is restarted.

Change the Speed of an Uplink Adapter

You can change the connection speed and duplex of an uplink adapter.

Procedure

- 1 Log in to the vSphere Client and select the host from the inventory panel.
- 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**.
- 6 To select the connection speed manually, select the speed and 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.

The adapter and the physical switch port it is connected to must be set to the same value, such as auto and auto or ND and ND, where ND is some speed and duplex, but not auto and ND.
- 7 Click **OK**.

Add Uplink Adapters

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

Procedure

- 1 Log in to the vSphere Client and select the host from the inventory panel.
- 2 Click the **Configuration** tab and click **Networking**.
- 3 Select a vSwitch and click **Properties**.
- 4 Click the **Network Adapters** tab.
- 5 Click **Add** to launch the Add Adapter wizard.
- 6 Select one or more adapters from the list and click **Next**.
- 7 (Optional) To reorder the NICs into a different category, select a NIC and click **Move Up** and **Move Down**.

Option	Description
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 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 vSphere Client.

Enable CDP on an ESX Host

vSwitches are set to detect Cisco port information by default. You can also set the CDP mode so that a vSwitch makes information available to the Cisco switch administrator.

Procedure

- 1 Log in directly to your ESX host's console.
- 2 View the current CDP mode for the a vSwitch by entering the **esxcfg-vswitch -b <vSwitch>** command.
If CDP is disabled, the mode will be shown as **down**.
- 3 Change the CDP mode by entering the **esxcfg-vswitch -B <mode> <vSwitch>** command.

Mode	Description
down	CDP is disabled.
listen	ESX 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 makes information about the vSwitch available to the Cisco switch administrator, but does not detect and display information about the Cisco switch.
both	ESX detects and displays information about the associated Cisco switch and makes information about the vSwitch available to the Cisco switch administrator.

View Cisco Switch Information on the vSphere Client

When CDP is set to **listen** or **both**, you can view Cisco switch information.

Procedure

- 1 Log in to the vSphere Client and select the host from the inventory panel.
- 2 Click the **Configuration** tab and click **Networking**.
- 3 Click the info icon to the right of the vSwitch.

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

Basic Networking with vNetwork Distributed Switches

4

These topics guide you through the basic concepts of networking with vNetwork Distributed Switches and how to set up and configure networking with vNetwork Distributed Switches in a vSphere environment.

This chapter includes the following topics:

- [“vNetwork Distributed Switch Architecture,”](#) on page 29
- [“Configuring a vNetwork Distributed Switch,”](#) on page 30
- [“dvPort Groups,”](#) on page 32
- [“Private VLANs,”](#) on page 34
- [“Configuring vNetwork Distributed Switch Network Adapters,”](#) on page 35
- [“Configuring Virtual Machine Networking on a vNetwork Distributed Switch,”](#) on page 39

vNetwork Distributed Switch Architecture

A vNetwork Distributed Switch functions as a single virtual switch across all associated hosts. This allows virtual machines to maintain a consistent network configuration as they migrate across multiple hosts.

Like a vNetwork Standard Switch, each vNetwork Distributed Switch is a network hub that virtual machines can use. A vNetwork Distributed Switch can forward traffic internally between virtual machines or link to an external network by connecting to physical Ethernet adapters, also known as uplink adapters.

Each vNetwork Distributed Switch can also have one or more dvPort groups assigned to it. dvPort groups group multiple ports under a common configuration and provide a stable anchor point for virtual machines connecting to labeled networks. Each dvPort group is identified by a network label, which is unique to the current datacenter. A VLAN ID, which restricts port group traffic to a logical Ethernet segment within the physical network, is optional.

In addition to VMware vNetwork Distributed Switches, vSphere 4 also provides initial support for third-party virtual switches. For information on configuring these third-party switches, visit <http://www.cisco.com/go/1000vdocs>.

Configuring a vNetwork Distributed Switch

You can create a vNetwork Distributed Switch on a vCenter Server datacenter. After you have created a vNetwork Distributed Switch, you can add hosts, create dvPort groups, and edit vNetwork Distributed Switch properties and policies.

Create a vNetwork Distributed Switch

Create a vNetwork Distributed Switch to handle networking traffic for associated hosts on the datacenter.

Procedure

- 1 Log in to the vSphere Client and display the datacenter in Networking view.
- 2 From the Inventory menu, select **Datacenter > vNetwork Distributed Switch**.
The Create vNetwork Distributed Switch wizard appears.
- 3 Enter a name for the vNetwork Distributed Switch in the Name field.
- 4 Select the **Number of dvUplink Ports**, and click **Next**.
dvUplink ports connect the vNetwork Distributed Switch to physical NICs on associated ESX hosts. The number of dvUplink ports is the maximum number of allowed physical connections to the vNetwork Distributed Switch per host.
- 5 Click **Next**.
- 6 Choose **Add now** or **Add later**.
- 7 If you chose **Add now**, select the hosts and physical adapters to use by clicking the check box next to each host or adapter. You can add only physical adapters that are not already in use during vNetwork Distributed Switch creation.
- 8 Click **Next**.
- 9 Choose whether to **Automatically create a default port group**.
This option creates an early-binding port group with 128 ports. For systems with complex port group requirements, skip the default port group and create a new dvPort group after you have finished adding the vNetwork Distributed Switch.
- 10 Review the vNetwork Distributed Switch diagram to ensure proper configuration, and click **Finish**.

What to do next

If you chose to add hosts later, you must add hosts to the vNetwork Distributed Switch before adding network adapters.

Network adapters can be added from the host configuration page of the vSphere Client or by using Host Profiles.

Add a Host to a vNetwork Distributed Switch

Use the Add Host to vNetwork Distributed Switch wizard to associate a host with a vNetwork Distributed Switch. You can also add hosts to a vNetwork Distributed Switch using Host Profiles.

Procedure

- 1 In the vSphere Client, display the Networking inventory view and select the vNetwork Distributed Switch.
- 2 From the Inventory menu, select **Distributed Virtual Switch > Add Host**.
The Add Host to vNetwork Distributed Switch wizard appears.

- 3 Select the host to add.
- 4 Under the selected host, select the physical adapters to add, and click **Next**.

You may select both free and in use physical adapters. If you select an adapter that is currently in use by a host, choose whether to move the associated virtual adapters to the vNetwork Distributed Switch.

NOTE Moving a physical adapter to a vNetwork Distributed Switch without moving any associated virtual adapters will cause those virtual adapters to lose network connectivity.

- 5 Click **Finish**.

Edit General vNetwork Distributed Switch Settings

You can edit the general properties for a vNetwork Distributed Switch, such as the vNetwork Distributed Switch name and the number of uplink ports on the vNetwork Distributed Switch.

Procedure

- 1 In the vSphere Client, display the Networking inventory view and select the vNetwork Distributed Switch.
- 2 From the Inventory menu, select **Distributed Virtual Switch > Edit Settings**.
- 3 Select **General** to edit the following vNetwork Distributed Switch settings.
 - a Enter the name for the vNetwork Distributed Switch.
 - b Select the number of uplink ports.
 - c To edit uplink port names, click **Edit uplink port names**, enter the new names, and click **OK**.
 - d Enter any notes for the vNetwork Distributed Switch.
- 4 Click **OK**.

Edit Advanced vNetwork Distributed Switch Settings

Use the vNetwork Distributed Switch Settings dialog box to configure advanced vNetwork Distributed Switch settings such as Cisco Discovery Protocol and the maximum MTU for the vNetwork Distributed Switch.

Procedure

- 1 In the vSphere Client, display the Networking inventory view and select the vNetwork Distributed Switch.
- 2 From the Inventory menu, select **Distributed Virtual Switch > Edit Settings**.
- 3 Select **Advanced** to edit the following vNetwork Distributed Switch properties.
 - a Specify the maximum MTU size.
 - b Select the **Enable Cisco Discovery Protocol** check box to enable CDP, and set the operation to **Listen**, **Advertise**, or **Both**.
 - c Enter the name and other details for the vNetwork Distributed Switch administrator in the Admin Contact Info section.
- 4 Click **OK**.

View Network Adapter Information for a vNetwork Distributed Switch

View physical network adapters and uplink assignments for a vNetwork Distributed Switch from the networking inventory view of the vSphere Client.

Procedure

- 1 In the vSphere Client, display the Networking inventory view and select the vNetwork Distributed Switch.

- 2 From the Inventory menu, select **Distributed Virtual Switch > Edit Settings**.
- 3 On the **Network Adapters** tab, you can view network adapter and uplink assignments for associated hosts.
This tab is read-only. vNetwork Distributed Switch network adapters must be configured at the host level.
- 4 Click **OK**.

dvPort Groups

A dvPort group specifies port configuration options for each member port on a vNetwork Distributed Switch. dvPort groups define how a connection is made to a network.

Add a dvPort Group

Use the Create dvPort Group wizard to add a dvPort group to a vNetwork Distributed Switch.

Procedure

- 1 In the vSphere Client, display the Networking inventory view and select the vNetwork Distributed Switch.
- 2 From the **Inventory** menu, select **Distributed Virtual Switch > New Port Group**.
- 3 Enter a name and the number of ports for the dvPort group.
- 4 Choose a VLAN type.

Option	Description
None	Do not use VLAN.
VLAN	In the VLAN ID field, enter a number between 1 and 4094.
VLAN Trunking	Enter a VLAN trunk range.
Private VLAN	Select a private VLAN entry. If you have not created any private VLANs, this menu is empty.

- 5 Click **Next**.
- 6 Click **Finish**.

Edit General dvPort Group Properties

Use the dvPort Group Properties dialog box to configure general dvPort group properties such as the dvPort group name and port group type.

Procedure

- 1 In the vSphere Client, display the Networking inventory view and select the dvPort group.
- 2 From the Inventory menu, select **Network > Edit Settings**.
- 3 Select **General** to edit the following dvPort group properties.

Option	Action
Name	Enter the name for the dvPort group.
Description	Enter a brief description of the dvPort group.

Option	Action
Number of Ports	Enter the number of ports on the dvPort group.
Port binding	Choose when ports are assigned to virtual machines connected to this dvPort group. ■ Select Static binding to assign a port to a virtual machine when the virtual machine is connected to the dvPort group. ■ Select Dynamic binding to assign a port to a virtual machine the first time the virtual machine powers on after it is connected to the dvPort group. ■ Select Ephemeral for no port binding.

- 4 Click **OK**.

Edit Advanced dvPort Group Properties

Use the dvPort Group Properties dialog box to configure advanced dvPort group properties such as the port name format and override settings.

Procedure

- 1 In the vSphere Client, display the Networking inventory view and select the dvPort group.
- 2 From the Inventory menu, select **Network > Edit Settings**.
- 3 Select **Advanced** to edit the dvPort group properties.
 - a Select **Allow port policies to be overridden** to allow dvPort group policies to be overridden on a per-port level.
 - b Click **Edit Override Settings** to select which policies can be overridden.
 - c Choose whether to allow live port moving.
 - d Select **Config reset at disconnect** to discard per-port configurations when a dvPort is disconnected from a virtual machine.
 - e Select **Binding on host allowed** to specify that when vCenter Server is down, ESX can assign a dvPort to a virtual machine.
 - f Select **Port name format** to provide a template for assigning names to the dvPorts in this group.
- 4 Click **OK**.

Configure dvPort Settings

Use the Port Settings dialog box to configure general dvPort properties such as the port name and description.

Procedure

- 1 Log in to the vSphere Client and display the vNetwork Distributed Switch.
- 2 On the **Ports** tab, right-click the port to modify and select **Edit Settings**.
- 3 Click **General**.
- 4 Modify the port name and description.
- 5 Click **OK**.

Private VLANs

Private VLANs are used to solve VLAN ID limitations and the waste of IP addresses for certain network setups.

A private VLAN is identified by its primary VLAN ID. A primary VLAN ID can have multiple secondary VLAN IDs associated with it. Primary VLANs are promiscuous so that ports on a private VLAN can communicate with ports configured as the primary VLAN. Ports on a secondary VLAN can be either isolated, communicating only with promiscuous ports, or community, communicating with both promiscuous ports and other ports on the same secondary VLAN.

To use private VLANs between an ESX host and the rest of the physical network, the physical switch connected to the ESX host needs to be private VLAN-capable and configured with the VLAN IDs being used by ESX for the private VLAN functionality. For physical switches using dynamic MAC+VLAN ID based learning, all corresponding private VLAN IDs must be first entered into the switch's VLAN database.

To configure dvPorts to use private VLAN functionality, you must first create the necessary private VLANs on the vNetwork Distributed Switch that the dvPorts are connected to.

Create a Private VLAN

You can create a private VLAN for use on a vNetwork Distributed Switch and its associated dvPorts.

Procedure

- 1 In the vSphere Client, display the Networking inventory view and select the vNetwork Distributed Switch.
- 2 From the **Inventory** menu, select **vNetwork Distributed Switch > Edit Settings**.
- 3 Select the **Private VLAN** tab.
- 4 Under Primary Private VLAN ID, click **[Enter a Private VLAN ID here]**, and enter the number of the primary private VLAN.
- 5 Click anywhere in the dialog box, and then select the primary private VLAN that you just added.
The primary private VLAN you added appears under Secondary Private VLAN ID.
- 6 For each new secondary private VLAN, click **[Enter a Private VLAN ID here]** under Secondary Private VLAN ID, and enter the number of the secondary private VLAN.
- 7 Click anywhere in the dialog box, select the secondary private VLAN that you just added, and select either **Isolated** or **Community** for the port type.
- 8 Click **OK**.

Remove a Primary Private VLAN

Remove unused primary private VLANs from the networking inventory view of the vSphere Client.

Prerequisites

Before removing a private VLAN, be sure that no port groups are configured to use it.

Procedure

- 1 In the vSphere Client, display the Networking inventory view and select the vNetwork Distributed Switch.
- 2 From the **Inventory** menu, select **vNetwork Distributed Switch > Edit Settings**.
- 3 Select the **Private VLAN** tab.
- 4 Select the primary private VLAN to remove.

- 5 Click **Remove** under Primary Private VLAN ID, and click **OK**.

Removing a primary private VLAN also removes all associated secondary private VLANs.

Remove a Secondary Private VLAN

Remove unused secondary private VLANs from the networking inventory view of the vSphere Client.

Prerequisites

Before removing a private VLAN, be sure that no port groups are configured to use it.

Procedure

- 1 In the vSphere Client, display the Networking inventory view and select the vNetwork Distributed Switch.
- 2 From the **Inventory** menu, select **vNetwork Distributed Switch > Edit Settings**.
- 3 Select the **Private VLAN** tab.
- 4 Select a primary private VLAN to display its associated secondary private VLANs.
- 5 Select the secondary private VLAN to remove.
- 6 Click **Remove** under Secondary Private VLAN ID, and click **OK**.

Configuring vNetwork Distributed Switch Network Adapters

The vNetwork Distributed Switch networking view of the host configuration page displays the configuration of the host's associated vNetwork Distributed Switches and allows you to configure the vNetwork Distributed Switch network adapters and uplink ports.

Managing Physical Adapters

For each host associated with a vNetwork Distributed Switch, you must assign physical network adapters, or uplinks, to the vNetwork Distributed Switch. You can assign one uplink on each host per uplink port on the vNetwork Distributed Switch.

Add an Uplink to a vNetwork Distributed Switch

Physical uplinks must be added to a vNetwork Distributed Switch in order for virtual machines and virtual network adapters connected to the vNetwork Distributed Switch to connect to networks outside the hosts on which they reside.

Procedure

- 1 Log in to the vSphere Client and select a host from the inventory panel.
- 2 Click the **Configuration** tab and click **Networking**.
- 3 Select the vNetwork Distributed Switch view.
- 4 Click **Manage Physical Adapters**.
- 5 Click **Click to Add NIC** for the uplink port to add an uplink to.
- 6 Select the physical adapter to add. If you select an adapter that is attached to another switch, it is removed from that switch and reassigned to this vNetwork Distributed Switch.
- 7 Click **OK**.

Remove an Uplink from a vNetwork Distributed Switch

An uplink that is associated with a vNetwork Distributed Switch cannot be added to a vSwitch or another vNetwork Distributed Switch.

Procedure

- 1 Log in to the vSphere Client and select a host from the inventory panel.
- 2 Click the **Configuration** tab and click **Networking**.
- 3 Select the vNetwork Distributed Switch view.
- 4 Click **Manage Physical Adapters**.
- 5 Click **Remove** for the uplink to remove.
- 6 Click **OK**.

Managing Virtual Network Adapters

Virtual network adapters handle host network services over a vNetwork Distributed Switch.

You can configure service console and VMkernel virtual adapters for an ESX host through an associated vNetwork Distributed Switch either by creating new virtual adapters or migrating existing virtual adapters.

Create a VMkernel Network Adapter on a vNetwork Distributed Switch

Create a VMkernel network adapter for use as a VMotion interface or an IP storage port group.

Procedure

- 1 Log in to the vSphere Client and select the host from the inventory panel.
- 2 Click the **Configuration** tab and click **Networking**.
- 3 Select the vNetwork Distributed Switch view.
- 4 Click **Manage Virtual Adapters**.
- 5 Click **Add**.
- 6 Select **New virtual adapter**, and click **Next**.
- 7 Select **VMkernel**, and click **Next**.
- 8 Under Network Connection, select the vNetwork Distributed Switch and the associated port group, or select the standalone port to which to add this virtual adapter.
- 9 Select **Use this virtual adapter for VMotion** to enable this port group to advertise itself to another ESX host as the network connection where VMotion traffic is sent.

You can enable this property for only one VMotion and IP storage port group for each ESX host. If this property is not enabled for any port group, migration with VMotion to this host is not possible.
- 10 Choose whether to use this virtual adapter for fault tolerance logging.
- 11 Under IP Settings, specify the IP address and subnet mask.
- 12 Click **Edit** to set the VMkernel default gateway for VMkernel services, such as VMotion, NAS, and iSCSI.

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

- 13 On the **Routing** tab, the service console and the VMkernel each need their own gateway information.
A gateway is needed for connectivity to machines not on the same IP subnet as the service console or VMkernel. Static IP settings is the default.
- 14 Click **OK**, and then click **Next**.
- 15 Click **Finish**.

Create a Service Console Network Adapter on a vNetwork Distributed Switch

You create a service console network adapter on a vNetwork Distributed Switch from the configuration page of an associated host.

Procedure

- 1 Log in to the vSphere Client and select the host from the inventory panel.
- 2 Click the **Configuration** tab and click **Networking**.
- 3 Select the vNetwork Distributed Switch view.
- 4 Click **Manage Virtual Adapters**.
- 5 Click **Add**.
- 6 Select **New virtual adapter**, and click **Next**.
- 7 Select **Service Console**, and click **Next**.
- 8 Under Network Connection, select the vNetwork Distributed Switch and the associated port group, or select the standalone port to which to add this virtual adapter.
- 9 Enter the IP address and subnet mask, or select **Obtain IP setting automatically**.
- 10 Click **Edit** to set the service console default gateway.
- 11 Click **Next**.
- 12 Click **Finish**.

Migrate an Existing Virtual Adapter to a vNetwork Distributed Switch

You migrate an existing virtual adapter from a vNetwork Standard Switch to a vNetwork Distributed Switch from the host configuration page.

Procedure

- 1 Log in to the vSphere Client and select the host from the inventory panel.
- 2 Click the **Configuration** tab and click **Networking**.
- 3 Select the vNetwork Distributed Switch view.
- 4 Click **Manage Virtual Adapters**.
- 5 Click **Add**.
- 6 Select **Migrate existing virtual network adapters**, and click **Next**.
- 7 In the **Select by** drop-down menu, choose whether to connect this virtual adapter to a port group or a standalone dvPort.
- 8 Select one or more virtual network adapters to migrate.
- 9 For each selected adapter, choose a port group or dvPort from the **Select a port group** or **Select a port** drop-down menu.
- 10 Click **Next**.

- 11 Click **Finish**.

Migrate a Virtual Adapter to a vNetwork Standard Switch

Use the Migrate to Virtual Switch wizard to migrate an existing virtual adapter from a vNetwork Distributed Switch to a vNetwork Standard Switch.

Procedure

- 1 Log in to the vSphere Client and select the host from the inventory panel.
The hardware configuration page for this server appears.
- 2 Click the **Configuration** tab and click **Networking**.
- 3 Select the **vNetwork Distributed Switch** view.
- 4 Click **Manage Virtual Adapters**.
- 5 Select the virtual adapter to migrate, and click **Migrate to Virtual Switch**.
The Migrate Virtual Adapter wizard appears.
- 6 Select the vSwitch to migrate the adapter to and click **Next**.
- 7 Enter a **Network Label** and optionally a **VLAN ID** for the virtual adapter, and click **Next**.
- 8 Click **Finish** to migrate the virtual adapter and complete the wizard.

Edit the VMkernel Configuration on a vNetwork Distributed Switch

You can edit the properties of an existing VMkernel adapter on a vNetwork Distributed Switch from the associated host.

Procedure

- 1 Log in to the vSphere Client and select the host from the inventory panel.
- 2 Click the **Configuration** tab and click **Networking**.
- 3 Select the vNetwork Distributed Switch view.
- 4 Click **Manage Virtual Adapters**.
- 5 Select the VMkernel adapter to modify and click **Edit**.
- 6 Under Network Connection, select the vNetwork Distributed Switch and the associated port group, or select the standalone port to which to add this virtual adapter.
- 7 Select **Use this virtual adapter for VMotion** to enable this port group to advertise itself to another ESX host as the network connection where VMotion traffic is sent.
You can enable this property for only one VMotion and IP storage port group for each ESX host. If this property is not enabled for any port group, migration with VMotion to this host is not possible.
- 8 Choose whether to use this virtual adapter for fault tolerance logging.
- 9 Under IP Settings, specify the IP address and subnet mask, or select **Obtain IP settings automatically**.
- 10 Click **Edit** to set the VMkernel default gateway for VMkernel services, such as VMotion, NAS, and iSCSI.
On the **DNS Configuration** tab, the name of the host is entered by default. The DNS server addresses and domain that were specified during installation are also preselected.
- 11 Click **OK**.

Edit the Service Console Configuration on a vNetwork Distributed Switch

You edit the properties of a service console virtual adapter from the networking view of the host configuration page.

Procedure

- 1 Log in to the vSphere Client and select the host from the inventory panel.
- 2 Click the **Configuration** tab and click **Networking**.
- 3 Select the vNetwork Distributed Switch view.
- 4 Click **Manage Virtual Adapters**.
- 5 Select the service console adapter to modify and click **Edit**.
- 6 Under Network Connection, select the vNetwork Distributed Switch and the associated port group, or select the standalone port to which to add this virtual adapter.
- 7 Enter the IP address and subnet mask, or select **Obtain IP setting automatically**.
- 8 Click **Edit** to set the service console default gateway.
- 9 Click **OK**.

Remove a Virtual Adapter

Remove a virtual network adapter from a vNetwork Distributed Switch in the Manage Virtual Adapters dialog box.

Procedure

- 1 Log in to the vSphere Client and select the host from the inventory panel.
- 2 Click the **Configuration** tab and click **Networking**.
- 3 Select the **vNetwork Distributed Switch** view.
- 4 Click **Manage Virtual Adapters**.
- 5 Select the virtual adapter to remove and click **Remove**.

A dialog box appears with the message, Are you sure you want to remove <adapter name>?

- 6 Click **Yes**.

Configuring Virtual Machine Networking on a vNetwork Distributed Switch

Virtual machines can be connected to a vNetwork Distributed Switch either by configuring an individual virtual machine NIC or migrating groups of virtual machines from the vNetwork Distributed Switch itself.

Virtual machines are connected to vNetwork Distributed Switches by connecting their associated virtual network adapters to dvPort groups. This can be done either for an individual virtual machine by modifying the virtual machine's network adapter configuration, or for a group of virtual machines by migrating virtual machines from an existing virtual network to a vNetwork Distributed Switch.

Migrate Virtual Machines to or from a vNetwork Distributed Switch

In addition to connecting virtual machines to a vNetwork Distributed Switch at the individual virtual machine level, you can migrate a group of virtual machines between a vNetwork Distributed Switch network and a vNetwork Standard Switch network.

Procedure

- 1 In the vSphere Client, display the Networking inventory view and select the vNetwork Distributed Switch.
- 2 From the **Inventory** menu, select **Distributed Virtual Switch > Migrate Virtual Machine Networking**.
The Migrate Virtual Machine Networking wizard appears.
- 3 In the **Select Source Network** drop-down menu, select the virtual network to migrate from.
- 4 Select the virtual network to migrate to from the **Select Destination Network** drop-down menu.
- 5 Click **Show Virtual Machines**.
Virtual machines associated with the virtual network you are migrating from are displayed in the **Select Virtual Machines** field.
- 6 Select virtual machines to migrate to the destination virtual network, and click **OK**.

Connect an Individual Virtual Machine to a dvPort Group

Connect an individual virtual machine to a vNetwork Distributed Switch by modifying the virtual machine's NIC configuration.

Procedure

- 1 Log in to the vSphere Client and select the virtual machine from the inventory panel.
- 2 On the **Summary** tab, click **Edit Settings**.
- 3 On the **Hardware** tab, select the virtual network adapter.
- 4 Select the dvPort group to migrate to from the **Network Label** drop-down menu, and click **OK**.

Advanced Networking

The following topics guide you through advanced networking in an ESX environment, and how to set up and change advanced networking configuration options.

This chapter includes the following topics:

- [“Internet Protocol Version 6,”](#) on page 41
- [“Networking Policies,”](#) on page 42
- [“Change the DNS and Routing Configuration,”](#) on page 55
- [“MAC Addresses,”](#) on page 55
- [“TCP Segmentation Offload and Jumbo Frames,”](#) on page 57
- [“NetQueue and Networking Performance,”](#) on page 60
- [“VMDirectPath Gen I,”](#) on page 60

Internet Protocol Version 6

vSphere supports both Internet Protocol version 4 (IPv4) and Internet Protocol version 6 (IPv6) environments.

The Internet Engineering Task Force has designated IPv6 as the successor to IPv4. The adoption of IPv6, both as a standalone protocol and in a mixed environment with IPv4, is rapidly increasing. With IPv6, you can use vSphere features such as NFS in an IPv6 environment.

A major difference between IPv4 and IPv6 is address length. IPv6 uses a 128-bit address rather than the 32-bit addresses used by IPv4. This helps alleviate the problem of address exhaustion that is present with IPv4 and eliminates the need for network address translation (NAT). Other notable differences include link-local addresses that appear as the interface is initialized, addresses that are set by router advertisements, and the ability to have multiple IPv6 addresses on an interface.

An IPv6-specific configuration in vSphere involves providing IPv6 addresses, either by entering static addresses or by using DHCP for all relevant vSphere networking interfaces. IPv6 addresses can also be configured using stateless autoconfiguration sent by router advertisement.

Enable IPv6 Support on an ESX Host

You can enable or disable IPv6 support on the host.

Procedure

- 1 Click the arrow next to the **Inventory** button in the navigation bar and select **Hosts and Clusters**.
- 2 Select the host and click the **Configuration** tab.
- 3 Click the **Networking** link under Hardware.

- 4 In the Virtual Switch view, click the **Properties** link.
- 5 Select **Enable IPv6 support on this host** and click **OK**.
- 6 Reboot the host.

Networking Policies

Any policies set at the vSwitch or dvPort group level are applied to all port groups on that vSwitch or dvPorts in the dvPort group, except for the configuration options that are overridden at the port group or dvPort level.

The following networking policies can be applied:

- Load balancing and failover
- VLAN (vNetwork Distributed Switch only)
- Security
- Traffic shaping
- Port blocking policies (vNetwork Distributed Switch only)

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 adapter failure.

You can edit your load balancing and failover policy 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** controls the link status and beacon probing. Beaconing is not supported with guest VLAN tagging.
- **Network Adapter Order** can be active or standby.

Edit the Failover and Load Balancing Policy on a vSwitch

Failover and load balancing 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..

Procedure

- 1 Log in to the vSphere Client and select the host from the inventory panel.
- 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 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 Specify the settings in the Policy Exceptions group.

Option	Description
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 uplinks. If you want to use some uplinks but reserve others for emergencies in case the uplinks in use fail, set this condition by moving them into different groups: ■ Active Uplinks — Continue to use the uplink when the network adapter connectivity is up and active. ■ Standby Uplinks — Use this uplink if one of the active adapter's connectivity is down. ■ Unused Uplinks — Do not use this uplink.

8 Click **OK**.

Edit the Failover and Load Balancing Policy on a Port Group

You can edit the failover and load balancing policy configuration for a port group.

Procedure

- 1 Log in to the vSphere Client and select the host from the inventory panel.
- 2 Click the **Configuration** tab, and click **Networking**.
- 3 Select a port group and click **Edit**.
- 4 In the 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 Specify the settings in the Policy Exceptions group.

Option	Description
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.

Option	Description
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 uplinks. If you want to use some uplinks but reserve others for emergencies in case the uplinks in use fail, set this condition by moving them into different groups: ■ Active Uplinks — Continue to use the uplink when the network adapter connectivity is up and active. ■ Standby Uplinks — Use this uplink if one of the active adapter's connectivity is down. ■ Unused Uplinks — Do not use this uplink.

- 8 Click **OK**.

Edit the Teaming and Failover Policy on a dvPort Group

Teaming 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.

Procedure

- 1 In the vSphere Client, display the Networking inventory view and select the dvPort group.
- 2 From the Inventory menu, select **Network > Edit Settings**.
- 3 Select **Policies**.
- 4 In the Teaming and Failover group, specify the following.

Option	Description
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. NOTE Do not use beacon probing with IP-has load balancing.

Option	Description
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 uplinks. If you want to use some uplinks but reserve others for emergencies in case the uplinks in use fail, set this condition by moving them into different groups: ■ Active Uplinks — Continue to use the uplink when the network adapter connectivity is up and active. ■ Standby Uplinks — Use this uplink if one of the active adapter's connectivity is down. ■ Unused Uplinks — Do not use this uplink. NOTE When using IP-has load balancing, do not configure standby uplinks.

- 5 Click **OK**.

Edit dvPort Teaming and Failover Policies

Teaming 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.

Prerequisites

To edit the teaming and failover policies on an individual dvPort, the associated dvPort group must be set to allow policy overrides.

Procedure

- 1 Log in to the vSphere Client and display the vNetwork Distributed Switch.
- 2 On the **Ports** tab, right-click the port to modify and select **Edit Settings**.
The **Port Settings** dialog box appears.
- 3 Click **Policies** to view and modify port networking policies.

- 4 In the Teaming and Failover group, specify the following.

Option	Description
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. NOTE Do not use beacon probing with IP-has load balancing.
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 uplinks. If you want to use some uplinks but reserve others for emergencies in case the uplinks in use fail, set this condition by moving them into different groups: ■ Active Uplinks — Continue to use the uplink when the network adapter connectivity is up and active. ■ Standby Uplinks — Use this uplink if one of the active adapter's connectivity is down. ■ Unused Uplinks — Do not use this uplink. NOTE When using IP-has load balancing, do not configure standby uplinks.

- 5 Click **OK**.

VLAN Policy

The VLAN policy allows virtual networks to join physical VLANs.

Edit the VLAN Policy on a dvPort Group

You can edit the VLAN policy configuration on a dvPort group.

Procedure

- 1 In the vSphere Client, display the Networking inventory view and select the dvPort group.
- 2 From the Inventory menu, select **Network > Edit Settings**.
- 3 Select **VLAN**.
- 4 Select the **VLAN** Type to use.

Option	Description
None	Do not use VLAN.
VLAN	In the VLAN ID field, enter a number between 1 and 4094.
VLAN Trunking	Enter a VLAN trunk range .
Private VLAN	Select an available private VLAN to use.

Edit dvPort VLAN Policies

A VLAN policy set at the dvPort level allows the individual dvPort to override the VLAN policy set at the dvPort group level.

Prerequisites

To edit the VLAN policies on an individual dvPort, the associated dvPort group must be set to allow policy overrides.

Procedure

- 1 Log in to the vSphere Client and display the vNetwork Distributed Switch.
- 2 On the **Ports** tab, right-click the port to modify and select **Edit Settings**.
- 3 Click **Policies**.
- 4 Select the VLAN type to use.

Option	Action
None	Do not use a VLAN.
VLAN	For the VLAN ID, enter a number between 1 and 4095.
VLAN Trunking	Enter a VLAN trunk range.
Private VLAN	Select an available private VLAN to use.

- 5 Click **OK**.

Security Policy

Networking security policies determine how the adapter filters inbound and outbound frames.

Layer 2 is the Data Link Layer. The three elements of the security policy are promiscuous mode, MAC address changes, and forged transmits.

In nonpromiscuous mode, a guest adapter listens only to traffic forwarded to own MAC address. In promiscuous mode, it can listen to all the frames. By default, guest adapters are set to nonpromiscuous mode.

Edit the Layer 2 Security Policy on a vSwitch

You can control how inbound and outbound frames are handled by editing Layer 2 security policies.

Procedure

- 1 Log in to the VMware vSphere Client and select the host from the inventory panel.
- 2 Click the **Configuration** tab, and click **Networking**.
- 3 Click **Properties** for the vSwitch to edit.
- 4 In the Properties dialog box, click the **Ports** tab.
- 5 Select the vSwitch item and click **Edit**.
- 6 In the Properties dialog box, click the **Security** tab.

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

The policy applies to all virtual adapters on the vSwitch, unless the port group for the virtual adapter specifies a policy exception.

- 7 In the Policy Exceptions pane, select whether to reject or accept the security policy exceptions.

Option	Description
Promiscuous Mode - Reject	Has no effect on which frames are received by the guest adapter.
Promiscuous Mode - Accept	The guest adapter detects 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 the guest OS 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 sent again.
MAC Address Changes - Accept	If the MAC address from the guest OS changes, frames to the new MAC address are received.
Forged Transmits -Reject	Outbound frames with a source MAC address that is different from the one set on the adapter are dropped.
Forged Transmits - Accept	No filtering is performed, and all outbound frames are passed.

- 8 Click **OK**.

Edit the Layer 2 Security Policy Exception on a Port Group

Control how inbound and outbound frames are handled by editing Layer 2 Security policies.

Procedure

- 1 Log in to the VMware vSphere Client and select the host from the inventory panel.
- 2 Click the **Configuration** tab, and click **Networking**.
- 3 Click **Properties** for the port group to edit.
- 4 In the Properties dialog box, click the **Ports** tab.
- 5 Select the port group item and click **Edit**.

- 6 In the Properties dialog box for the port group, click the **Security** tab.

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

The policy exception overrides any policy set at the vSwitch level.

- 7 In the Policy Exceptions pane, select whether to reject or accept the security policy exceptions.

Option	Description
Promiscuous Mode - Reject	Has no effect on which frames are received by the guest adapter.
Promiscuous Mode - Accept	The guest adapter detects 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 the guest OS 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 sent again.
MAC Address Changes - Accept	If the MAC address from the guest OS changes, frames to the new MAC address are received.
Forged Transmits -Reject	Outbound frames with a source MAC address that is different from the one set on the adapter are dropped.
Forged Transmits - Accept	No filtering is performed, and all outbound frames are passed.

- 8 Click **OK**.

Edit the Security Policy on a dvPort Group

Control how inbound and outbound frames for a dvPort group are handled by editing the Security policies.

Procedure

- 1 In the vSphere Client, display the Networking inventory view and select the dvPort group.
- 2 From the Inventory menu, select **Network > Edit Settings**.
- 3 In the Properties dialog box for the port group, click the **Security** tab.

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

The policy exception overrides any policy set at the vSwitch level.

- 4 In the Policy Exceptions pane, select whether to reject or accept the security policy exceptions.

Option	Description
Promiscuous Mode - Reject	Has no effect on which frames are received by the guest adapter.
Promiscuous Mode - Accept	The guest adapter detects 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 the guest OS 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 sent again.
MAC Address Changes - Accept	If the MAC address from the guest OS changes, frames to the new MAC address are received.
Forged Transmits -Reject	Outbound frames with a source MAC address that is different from the one set on the adapter are dropped.
Forged Transmits - Accept	No filtering is performed, and all outbound frames are passed.

- 5 Click **OK**.

Edit dvPort Security Policies

Control how inbound and outbound frames for a dvPort are handled by editing the Security policies.

Prerequisites

To edit the Security policies on an individual dvPort, the associated dvPort group must be set to allow policy overrides.

Procedure

- 1 Log in to the vSphere Client and display the vNetwork Distributed Switch.
- 2 On the **Ports** tab, right-click the port to modify and select **Edit Settings**.
- 3 Click **Policies**.

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

- 4 In the Security group, select whether to reject or accept the security policy exceptions:

Option	Description
Promiscuous Mode - Reject	Has no effect on which frames are received by the guest adapter.
Promiscuous Mode - Accept	The guest adapter detects 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 the guest OS 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 sent again.
MAC Address Changes - Accept	If the MAC address from the guest OS changes, frames to the new MAC address are received.
Forged Transmits -Reject	Outbound frames with a source MAC address that is different from the one set on the adapter are dropped.
Forged Transmits - Accept	No filtering is performed, and all outbound frames are passed.

- 5 Click **OK**.

Traffic Shaping Policy

A traffic shaping policy is defined by three characteristics: average bandwidth, peak bandwidth, and burst size. You can establish a traffic shaping policy for each port group and each dvPort or dvPort group.

ESX shapes outbound network traffic on vSwitches and both inbound and outbound traffic on a vNetwork Distributed Switch. Traffic shaping restricts the network bandwidth available on a port, but can also be configured to allow “bursts” of traffic to flow through at higher speeds.

Average Bandwidth	Establishes the number of bits per second to allow across a port, averaged over time—the allowed average load.
Peak Bandwidth	The maximum number of bits per second to allow across a port when it is sending or receiving a burst of traffic. This tops the bandwidth used by a port whenever it is using its burst bonus.
Burst Size	The maximum number of bytes to allow in a burst. If this parameter is set, a port might gain a burst bonus if it does not use all its allocated bandwidth. Whenever the port needs more bandwidth than specified by Average Bandwidth , it might be allowed to temporarily transmit data at a higher speed if a burst bonus is available. This parameter tops the number of bytes that have accumulated in the burst bonus and thus transfers at a higher speed.

Edit the Traffic Shaping Policy on a vSwitch

Use traffic shaping policies to control the bandwidth and burst size on a vSwitch.

Procedure

- 1 Log in to the vSphere Client and select the host from the inventory panel.
- 2 Click the **Configuration** tab, and click **Networking**.
- 3 Click **Properties** for the vSwitch to edit.
- 4 In the Properties dialog box, click the **Ports** tab.
- 5 Select the vSwitch item and click **Edit**.
- 6 In the Properties dialog box, click the **Traffic Shaping** tab.

When traffic shaping is disabled, the options 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.

NOTE Peak bandwidth cannot be less than the specified average bandwidth.

Option	Description
Status	If you enable the policy exception in the Status field, you are setting limits on the amount of networking bandwidth allocated for each virtual adapter associated with this particular port group. If you disable the policy, services have a free and clear connection to the physical network.
Average Bandwidth	A value measured over a particular period of time.
Peak Bandwidth	Limits the maximum bandwidth during a burst. It can never be smaller than the average bandwidth.
Burst Size	Specifies how large a burst can be in kilobytes (KB).

Edit the Traffic Shaping Policy on a Port Group

Use traffic shaping policies to control the bandwidth and burst size on a port group.

Procedure

- 1 Log in to the vSphere Client and select the host from the inventory panel.
- 2 Click the **Configuration** tab, and click **Networking**.
- 3 Click **Properties** for the port group to edit.
- 4 In the Properties dialog box, click the **Ports** tab.
- 5 Select the port group item and click **Edit**.
- 6 In the Properties dialog box for the port group, click the **Traffic Shaping** tab.

When traffic shaping is disabled, the options are dimmed.

Option	Description
Status	If you enable the policy exception in the Status field, you are setting limits on the amount of networking bandwidth allocated for each virtual adapter associated with this particular port group. If you disable the policy, services have a free and clear connection to the physical network.
Average Bandwidth	A value measured over a particular period of time.
Peak Bandwidth	Limits the maximum bandwidth during a burst. It can never be smaller than the average bandwidth.
Burst Size	Specifies how large a burst can be in kilobytes (KB).

Edit the Traffic Shaping Policy on a dvPort Group

You can shape both inbound and outbound traffic on vNetwork Distributed Switches. You can restrict the network bandwidth available to a port, but you can also temporarily allow bursts of traffic to flow through a port at higher speeds.

Procedure

- 1 In the vSphere Client, display the Networking inventory view and select the dvPort group.
- 2 From the Inventory menu, select **Network > Edit Settings**.
- 3 Select **Traffic Shaping**.
- 4 In the Properties dialog box for the port group, click the **Traffic Shaping** tab.

You can configure both inbound traffic shaping and outbound traffic shaping. When traffic shaping is disabled, the options are dimmed.

NOTE Peak bandwidth cannot be less than the specified average bandwidth.

Option	Description
Status	If you enable the policy exception in the Status field, you are setting limits on the amount of networking bandwidth allocated for each virtual adapter associated with this particular port group. If you disable the policy, services have a free and clear connection to the physical network.
Average Bandwidth	A value measured over a particular period of time.
Peak Bandwidth	Limits the maximum bandwidth during a burst. It can never be smaller than the average bandwidth.
Burst Size	Specifies how large a burst can be in kilobytes (KB).

Edit dvPort Traffic Shaping Policies

You can shape both inbound and outbound traffic on vNetwork Distributed Switches. You can restrict the network bandwidth available to a port, but you can also temporarily allow bursts of traffic to flow through a port at higher speeds.

A traffic shaping policy is defined by three characteristics: average bandwidth, peak bandwidth, and burst size.

Prerequisites

To edit the traffic shaping policies on an individual dvPort, the associated dvPort group must be set to allow policy overrides.

Procedure

- 1 Log in to the vSphere Client and display the vNetwork Distributed Switch.
- 2 On the **Ports** tab, right-click the port to modify, and select **Edit Settings**.
- 3 Click **Policies**.
- 4 In the Traffic Shaping group, you can configure both inbound traffic shaping and outbound traffic shaping. When traffic shaping is disabled, the options are dimmed.

Option	Description
Status	If you enable the policy exception in the Status field, you are setting limits on the amount of networking bandwidth allocated for each virtual adapter associated with this particular port group. If you disable the policy, services have a free and clear connection to the physical network.
Average Bandwidth	A value measured over a particular period of time.
Peak Bandwidth	Limits the maximum bandwidth during a burst. It can never be smaller than the average bandwidth.
Burst Size	Specifies how large a burst can be in kilobytes (KB).

- 5 Click **OK**.

Port Blocking Policies

Set blocking policies for dvPorts from the miscellaneous policies dialog box.

Edit the Port Blocking Policy on a dvPort Group

Set the port blocking policy for a dvPort group under miscellaneous policies.

Procedure

- 1 In the vSphere Client, display the Networking inventory view and select the dvPort group.
- 2 From the Inventory menu, select **Network > Edit Settings**.
- 3 Select **Miscellaneous**.
- 4 Choose whether to **Block all ports** on this dvPort group.

Edit dvPort Port Blocking Policy

The Miscellaneous policies dialog allows you to configure port blocking policies for a dvPort.

Procedure

- 1 Log in to the vSphere Client and display the vNetwork Distributed Switch.
- 2 On the **Ports** tab, right-click the port to modify and select **Edit Settings**.
- 3 Click **Policies**.
- 4 In the **Miscellaneous** group, select whether to **Block all ports**.
- 5 Click **OK**.

Change the DNS and Routing Configuration

You can change the DNS server and default gateway information provided during installation from the host configuration page in the vSphere Client.

Procedure

- 1 Log in to the vSphere Client and select the host from the inventory panel.
- 2 Click the **Configuration** tab, and click **DNS and Routing**.
- 3 On the right side of the window, click **Properties**.
- 4 In the **DNS Configuration** tab, enter a name and domain.
- 5 Choose whether 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. 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 the 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**.

MAC Addresses

MAC addresses are generated for virtual network adapters that the service console, the VMkernel, and virtual machines use.

In most cases, the generated 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 hosts 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.

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 the Organizationally Unique Identifier (OUI) 00:50:56 for manually generated addresses.

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 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 virtual machines do not support arbitrary MAC addresses, you must use the above format. 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.

MAC Address Generation

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

VMware has the following OUIs:

- Generated MAC addresses
- Manually set MAC addresses
- For legacy virtual machines, but no longer used with ESX

The first three bytes of the MAC address that is generated for each virtual network adapter consists of the OUI. The 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 host. 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 conflicts. The generated MAC addresses are created by using three parts: the VMware OUI, the SMBIOS UUID for the physical ESX machine, and a hash based on the name of the entity that the MAC address is being generated for.

After the MAC address has been 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 machine are tracked.

The MAC address of a powered off virtual machine is not checked against those of running or suspended virtual machines. It is possible that when a virtual machine is powered on again, it can acquire a different MAC address. This acquisition is caused by a conflict with a virtual machine that was powered on when this virtual machine was powered off.

Set Up a MAC Address

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

Procedure

- 1 Log in to the vSphere 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**.

TCP Segmentation Offload and Jumbo Frames

Jumbo frames must be enabled at the host level using the command-line interface to configure the MTU size for each vSwitch. TCP Segmentation Offload (TSO) is enabled on the VMkernel interface by default, but must be enabled at the virtual machine level.

Enabling TSO

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 might result in a change in the MAC address of the virtual network adapter.

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)

Enable TSO Support for a Virtual Machine

You can enable TSO support on a virtual machine by using an enhanced vmxnet adapter for that virtual machine.

Procedure

- 1 Log in to the vSphere 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 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** and then click **OK**.
- 11 If the virtual machine is not set to upgrade VMware Tools at each power on, you must upgrade VMware Tools manually.

TSO is enabled on a VMkernel interface. If TSO becomes disabled for a particular VMkernel interface, the only way to enable TSO is to delete that VMkernel interface and recreate it with TSO enabled.

Check Whether TSO Is Enabled on a VMkernel Interface

You can check whether TSO is enabled on a particular VMkernel networking interface.

Procedure

- 1 Log in to your ESX host's console.
- 2 Use the `esxcfg-vmknic -l` command to display a list of VMkernel interfaces.

The list shows each TSO-enabled VMkernel interface with TSO MSS set to 65535.

What to do next

If TSO is not enabled for a particular VMkernel interface, the only way to enable it is to delete the VMkernel interface and recreate the interface.

Enabling Jumbo Frames

Jumbo frames allow ESX to send larger frames out onto the physical network. The network must support jumbo frames end-to-end.

Jumbo frames up to 9kB (9000 bytes) are supported.

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

Create a Jumbo Frames-Enabled vSwitch

You configure a vSwitch for jumbo frames by changing the MTU size for that vSwitch.

Procedure

- 1 Use the `vicfg-vswitch -m <MTU> <vSwitch>` command in the VMware vSphere CLI to set the MTU size for the vSwitch.

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

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

Enable Jumbo Frames on a vNetwork Distributed Switch

You enable a vNetwork Distributed Switch for jumbo frames by changing the MTU size for that vNetwork Distributed Switch.

Procedure

- 1 In the vSphere Client, display the Networking inventory view and select the vNetwork Distributed Switch.
- 2 From the Inventory menu, select **Distributed Virtual Switch > Edit Settings**.
- 3 On the **Properties** tab, select **Advanced**.
- 4 Set the **Maximum MTU** to the largest MTU size among all the virtual network adapters connected to the vNetwork Distributed Switch, and click **OK**.

Enable Jumbo Frame Support on a Virtual Machine

Enabling jumbo frame support on a virtual machine requires an enhanced vmxnet adapter for that virtual machine.

Procedure

- 1 Log in to the vSphere 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 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 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.

Create a Jumbo Frames-Enabled VMkernel Interface

You can create a VMkernel network interface enabled with jumbo frames.

Procedure

- 1 Log in to your ESX 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 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-side networking performance.

Enable NetQueue on an ESX Host

NetQueue is enabled by default. To use NetQueue after it has been disabled, you must reenable it.

Procedure

- 1 Log in to the vSphere Client and select the host from the inventory panel.
- 2 Click the **Configuration** tab, and click **Advanced Settings** from the **Software** menu.
- 3 Select **VMkernel**.
- 4 Select **VMkernel.Boot.netNetQueueEnable** and click **OK**.
- 5 Use the VMware vSphere CLI to configure the NIC driver to use NetQueue.
See the *VMware vSphere Command-Line Interface Installation and Reference* guide.
- 6 Reboot the ESX host.

Disable NetQueue on an ESX Host

NetQueue is enabled by default.

Procedure

- 1 Log in to the vSphere Client and select the host from the inventory panel.
- 2 Click the **Configuration** tab, and click **Advanced Settings**.
- 3 Deselect **VMkernel.Boot.netNetQueueEnable** and click **OK**.
- 4 To disable NetQueue on the NIC driver, use the `vicfg-module -s "" [module name]` command.
For example, if you are using the s2io NIC driver, use `vicfg-module -s "" s2io`.
For information on the VMware vSphere CLI, see the *VMware vSphere Command-Line Interface Installation and Reference* guide.
- 5 Reboot the host.

VMDirectPath Gen I

With vSphere 4, ESX supports a direct PCI device connection for virtual machines running on Intel Nehalem platforms. Each virtual machine can connect up to two passthrough devices.

The following features are unavailable for virtual machines configured with VMDirectPath:

- VMotion
- Hot adding and removing of virtual devices
- Suspend and resume
- Record and replay
- Fault tolerance
- High availability

- DRS (limited availability; the virtual machine can be part of a cluster, but cannot migrate across hosts)

Configure Passthrough Devices on a Host

You can configure passthrough networking devices on a host.

Procedure

- 1 Select a host from the inventory panel of the vSphere Client.
- 2 On the **Configuration** tab, click **Advanced Settings**.

The Passthrough Configuration page appears, listing all available passthrough devices. A green icon indicates that a device is enabled and active. An orange icon indicates that the state of the device has changed and the host must be rebooted before the device can be used.

- 3 Click **Edit**.
- 4 Select the devices to be used for passthrough and click **OK**.

Configure a PCI Device on a Virtual Machine

You can configure a passthrough PCI device on a virtual machine.

Procedure

- 1 Select a virtual machine from the inventory panel of the vSphere Client.
- 2 From the **Inventory** menu, select **Virtual Machine > Edit Settings**.
- 3 On the **Hardware** tab, click **Add**.
- 4 Select **PCI Device** and click **Next**.
- 5 Select the passthrough device to use, and click **Next**.
- 6 Click **Finish**.

Adding a VMDirectPath device to a virtual machine sets memory reservation to the memory size of the virtual machine.

Networking Best Practices, Scenarios, and Troubleshooting

6

These topics describe networking best practices and common networking configuration and troubleshooting scenarios.

This chapter includes the following topics:

- [“Networking Best Practices,”](#) on page 63
- [“Mounting NFS Volumes,”](#) on page 64
- [“Networking Configuration for Software iSCSI Storage,”](#) on page 64
- [“Configuring Networking on Blade Servers,”](#) on page 65
- [“Troubleshooting,”](#) on page 67

Networking Best Practices

When configuring your network, consider these best practices.

- 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 separate physical networks (the latter is preferable).
 - Keeping the service console on its own network is an important part of securing the ESX system. Consider the service console network connectivity in the same light as any remote access device in a host, because compromising 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 is transmitted over the network.
- When using passthrough devices with a Linux kernel version 2.6.20 or earlier, avoid MSI and MSI-X modes because these modes have significant performance impact.
- 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 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 and that 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 the running hardware, the virtual machines can still communicate among themselves. Moreover, if you leave one NIC intact, all the virtual machines can still connect with the physical network.
- 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.

Mounting NFS Volumes

In ESX, the model of how ESX 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 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 vSphere 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 Configuration for Software iSCSI Storage

The storage you configure for an ESX 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 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.

Create a VMkernel Port for Software iSCSI

Before you can configure iSCSI storage, you must create one or more VMkernel ports to handle iSCSI networking.

Procedure

- 1 Log in to the vSphere Client and select the host from the inventory panel.
- 2 Click the **Configuration** tab, and click **Networking**.
- 3 Click **Add Networking**.
- 4 Select **VMkernel** and click **Next**.

On the Network Access 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 for the vSwitch.

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, existing vSwitches are using all the adapters.

Your choices appear in the Preview pane.

NOTE Do not use iSCSI on 100Mbps or slower network adapters.

- 7 Click **Next**.

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

Enter a network label to identify the port group that you are creating. When you configure iSCSI storage, specify this label.

Enter a VLAN ID to identify the VLAN that the port group's network traffic will use. VLAN IDs are not required. If you are not sure whether 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** and then click **Next**.
- 11 Click **Back** to make any changes.
- 12 Review your changes on the Ready to Complete page and click **Finish**.

Configuring Networking on Blade Servers

Because blade servers 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, 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 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 list describes how to configure a tagged port on three of the most commonly used blade servers.

Table 6-1. Configuration Options for Blade Servers

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

Configure a Virtual Machine Port Group with a VLAN on a Blade Server

Configuring virtual machine networking on a blade server requires some special considerations.

Procedure

- 1 Log in to the vSphere Client and select the host from the inventory panel.
- 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 **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 For **VLAN ID**, 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**.

Configure a VMkernel Port with a VLAN on a Blade Server

You can configure a VMkernel networking interface using a VLAN on a blade server.

Procedure

- 1 Log in to the vSphere Client and select the host from the inventory panel.
- 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.
Enter a network label to identify 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.
Enter a VLAN ID to identify 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 host 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 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.
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** and then click **Next**.
- 10 Click **Back** to make any changes.
- 11 Review your changes on the Ready to Complete page and click **Finish**.

Troubleshooting

The following topics guide you through troubleshooting common networking issues that you might encounter in an ESX environment.

Troubleshooting Service Console Networking

If certain parts of the service console's networking are misconfigured, you cannot access your ESX host with the vSphere Client.

If your host's service console loses network connectivity, you can reconfigure networking by connecting directly to the service console and using 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 the current virtual switch configurations. Check that the uplink adapter configured for the service console is connected to the appropriate physical network.

- `esxcfg-nics -l`

Provides a list of the 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, 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.

Rename Network Adapters by Using the Service Console

If you lose service console connectivity after adding a new network adapter, you must use the service console to rename the affected network adapters. Adding a new network adapter can cause loss of service console connectivity and manageability by using the vSphere Client because of network adapters are renamed.

Procedure

- 1 Log in directly to the ESX host's console.

- 2 Use the `esxcfg-nics -l` command to see which names are assigned to the network adapters.
- 3 Use the `esxcfg-vswitch -l` command to see which vSwitches 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 add the network adapters again, giving them the correct names.

Troubleshooting Physical Switch Configuration

You might lose vSwitch connectivity when a failover or failback event occurs. This causes the MAC addresses that the virtual machines associated with that vSwitch to appear on a different switch port.

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 an 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

This introduction describes available storage options for ESX and explains how to configure your ESX system so that it can use and manage different types of storage.

This chapter includes the following topics:

- [“About ESX Storage,”](#) on page 71
- [“Types of Physical Storage,”](#) on page 72
- [“Supported Storage Adapters,”](#) on page 73
- [“Target and Device Representations,”](#) on page 73
- [“About ESX Datastores,”](#) on page 76
- [“Comparing Types of Storage,”](#) on page 78
- [“Viewing Storage Information in the vSphere Client,”](#) on page 79

About ESX Storage

ESX storage is storage space on a variety of physical storage systems, local or networked, that a host uses to store virtual machine disks.

A 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, a host requires dedicated storage space.

The host uses storage space on a variety of physical storage systems, including your host’s internal and external devices, or networked storage, dedicated to the specific tasks of storing and protecting data.

The host can discover storage devices to which it has access and format them as datastores. The datastore is a special logical container, analogous to a file system on a logical volume, where ESX 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 vSphere Client, you can set up datastores on any storage device that your host discovers. In addition, you can use folders to create logical groups of datastores for organizational purposes, and for setting permissions and alarms across the datastore group.

Types of Physical Storage

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

ESX supports the following types of storage:

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

Local Storage

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

Local storage does not require a storage network to communicate with your host. All you need is a cable connected to the storage unit and, when required, a compatible HBA in your host.

Generally, you can connect multiple 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.

When multiple hosts connect to the local storage unit, they access storage devices in the unshared mode. The unshared mode does not permit several hosts to access the same VMFS datastore concurrently. However, a few SAS storage systems offer shared access to multiple hosts. This type of access permits multiple hosts to access the same VMFS datastore on a LUN.

ESX supports a variety of internal or external local storage devices, including SCSI, IDE, SATA, USB, and SAS storage systems. No matter which type of storage you use, your host 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 hosts.
- Some SAS storage systems can offer shared access to the same LUNs (and, therefore, the same VMFS datastores) to multiple hosts.

Networked Storage

Networked storage consists of external storage systems that your ESX host uses to store virtual machine files remotely. The host accesses these systems over a high-speed storage network.

ESX supports the following networked storage technologies.

NOTE Accessing the same storage through different transport protocols, such as iSCSI and Fibre Channel, at the same time is not supported.

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 hosts to high-performance storage devices. The network uses Fibre Channel protocol to transport SCSI traffic from virtual machines to the FC SAN devices.
---------------------------	--

Internet SCSI (iSCSI)

To connect to the FC SAN, your 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.

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 an iSCSI connection, your host serves as the initiator that communicates with a target, located in remote iSCSI storage systems.

ESX offers the following types of iSCSI connections:

Hardware-initiated iSCSI

Your host connects to storage through a third-party iSCSI HBA.

Software-initiated iSCSI

Your host 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 a standard TCP/IP network. The NFS client built into ESX uses Network File System (NFS) protocol version 3 to communicate with the NAS/NFS servers. For network connectivity, the host requires a standard network adapter.

Supported Storage Adapters

Storage adapters provide connectivity for your ESX host to a specific storage unit or network.


Depending on the type of storage you use, you might need to install or enable a storage adapter on your host. ESX supports different classes of adapters, including SCSI, iSCSI, RAID, Fibre Channel, and Ethernet. ESX accesses the adapters directly through device drivers in the VMkernel.

Target and Device Representations

In the ESX context, the term target identifies a single storage unit that the host can access. The terms device and LUN describe a logical volume that represents storage space on a target. Typically, the terms device and LUN, in the ESX context, mean a SCSI volume presented to the host from a storage target and available for formatting.

Different storage vendors present the storage systems to ESX hosts in different ways. Some vendors present a single target with multiple storage devices or LUNs on it, while others present multiple targets with one LUN each.

Figure 7-1. Target and LUN Representations



In this illustration, three LUNs are available in each configuration. In one case, the host sees one target, but that target has three LUNs that can be used. Each LUN represents an individual storage volume. In the other example, the host sees three different targets, each having one LUN.

Targets that are accessed through the network have unique names that are provided by the storage systems. The iSCSI targets use iSCSI names, while Fibre Channel targets use World Wide Names (WWNs).

NOTE ESX does not support accessing the same LUN through different transport protocols, such as iSCSI and Fibre Channel.

A device, or LUN, is identified by its UUID name.

Understanding Fibre Channel Naming

In Fibre Channel SAN, a World Wide Name (WWN) uniquely identifies each element in the network, such as a Fibre Channel adapter or storage device.

The WWN is a 64-bit address that consists of 16 hexadecimal numbers and might look like this:

20:00:00:e0:8b:8b:38:77 21:00:00:e0:8b:8b:38:77

The WWN is assigned to every Fibre Channel SAN element by its manufacturer.

Understanding iSCSI Naming and Addressing

In an iSCSI network, each iSCSI element that uses the network has a unique and permanent iSCSI name and is assigned an address for access.

iSCSI Name

Identifies a particular iSCSI element, regardless of its physical location. The iSCSI name can use IQN or EUI format.

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

`iqn.yyyy-mm.naming-authority:unique name`

- `yyyy-mm` is the year and month when the naming authority was established.
- `naming-authority` is usually reverse syntax of the Internet domain name of the naming authority. For example, the `iscsi.vmware.com` naming authority could have the iSCSI qualified name form of `iqn.1998-01.com.vmware.iscsi`. The name indicates that the `vmware.com` domain name was registered in January of 1998, and `iscsi` is a subdomain, maintained by `vmware.com`.
- `unique name` is any name you want to use, for example, the name of your host. The naming authority must make sure that any names assigned following the colon are unique, such as:
 - `iqn.1998-01.com.vmware.iscsi:name1`
 - `iqn.1998-01.com.vmware.iscsi:name2`
 - `iqn.1998-01.com.vmware.iscsi:name999`

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

For example,

```
eui.0123456789ABCDEF
```

iSCSI Alias

A more manageable, easy-to-remember name to use instead of the iSCSI name. iSCSI aliases are not unique, and are intended to be just a friendly name to associate with the node.

IP Address

An address associated with each iSCSI element so that routing and switching equipment on the network can establish the connection between different elements, such as the host and storage. This is just like the IP address you assign to a computer to get access to your company's network or the Internet.

Understanding Storage Device Naming

In the vSphere Client, each storage device, or LUN, is identified by several names, including a friendly name, a UUID, and a runtime name.

Name

A friendly name that the ESX host assigns to a device based on the storage type and manufacturer. You can modify the name using the vSphere Client. When you modify the name of the device on one host, the change takes affect across all hosts that have access to this device.

Identifier

A universally unique identifier assigned to a device. Depending on the type of storage, different algorithms are used to create the identifier. The identifier is persistent across reboots and is the same for all hosts sharing the device.

Runtime Name

The name of the first path to the device. The runtime name is created by the host, is not a reliable identifier for the device, and is not persistent.

The runtime name has the following format: `vmhba#:C#:T#:L#`, where

- `vmhba#` is the name of the storage adapter. The name refers to the physical adapter on the host, not to the SCSI controller used by the virtual machines.
- `C#` is the storage channel number.

Software iSCSI initiators use the channel number to show multiple paths to the same target.

- `T#` is the target number. Target numbering is decided by the host and might change if there is a change in the mappings of targets visible to the host. Targets that are shared by different ESX hosts might not have the same target number.
- `L#` is the LUN number that shows the position of the LUN within the target. The LUN number is provided by the storage system. If a target has only one LUN, the LUN number is always zero (0).

For example, `vmhba1:C0:T3:L1` represents LUN1 on target 3 accessed through the storage adapter `vmhba1` and channel 0.

About ESX Datastores

Datastores are 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.

You use the vSphere Client to access different types of storage devices that your ESX host discovers and to deploy datastores on them.

Depending on the type of storage you use, datastores can be backed by the following file system formats:

Virtual Machine File System (VMFS)

High-performance file system optimized for storing virtual machines. Your host can deploy a VMFS datastore 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.

Network File System (NFS)

File system on a NAS storage device. ESX supports NFS version 3 over TCP/IP. The host can access a designated NFS volume located on an NFS server, mount the volume, and use it for any storage needs.

If you use the service console to access your ESX host, you can see the VMFS and NFS datastores as separate subdirectories in the `/vmfs/volumes` directory.

VMFS Datastores

ESX can format SCSI-based storage devices 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 can set up VMFS datastores on any SCSI-based storage devices that your ESX host discovers. After you create the VMFS datastore, you can edit its properties.

You can have up to 256 VMFS datastores per system, with a minimum volume size of 1.2GB.

NOTE Always have only one VMFS datastore for each LUN.


If your VMFS datastore requires more space, you can increase the VMFS volume. You can dynamically add new extents to any VMFS datastore and grow the datastore up to 64TB. An extent is a LUN or partition on a physical storage device. The datastore can stretch over multiple extents, yet appear as a single volume.

Another option is to grow the existing datastore extent if the storage device where your datastore resides has free space. You can grow the extent up to 2 TB.

Sharing a VMFS Volume Across ESX Hosts

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

Figure 7-2. Sharing a VMFS Volume Across 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 hosts offers 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.
- You can use vMotion to migrate running virtual machines from one physical server to another.
- 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.

NFS Datastore

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

ESX supports the following shared storage capabilities on NFS volumes:

- vMotion
- VMware DRS and VMware HA
- ISO images, which are presented as CD-ROMs to virtual machines

- Virtual machine snapshots


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 host uses to connect to a storage device.

ESX supports Fibre Channel (FC), Internet SCSI (iSCSI), and NFS protocols. Regardless of the type of storage device your host 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 operating systems that are not certified for specific storage equipment, such as SAN, inside the virtual machine.

Figure 7-3 depicts five virtual machines using different types of storage to illustrate the differences between each type.

Figure 7-3. 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

Whether certain vSphere functionality is supported might depend on the storage technology that you use.

Table 7-1 compares networked storage technologies that ESX supports.

Table 7-1. Networked Storage that ESX 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 7-2 compares the vSphere features that different types of storage support.

Table 7-2. vSphere Features Supported by Storage

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

See *Setup for Failover Clustering and Microsoft Cluster Service* for support information concerning Microsoft clustering.

Viewing Storage Information in the vSphere Client

The vSphere Client displays detailed information on storage adapters and devices, and any available datastores.

Displaying Storage Adapters

The host uses storage adapters to access different storage devices. You can display the available storage adapters and review their information.

[Table 7-3](#) lists information that you can see when you display details for each adapter. Certain adapters, for example iSCSI, need to be configured or enabled before you can view their information.

Table 7-3. Storage Adapter Information

Adapter Information	Description
Model	Model of the adapter.
Targets (Fibre Channel and SCSI)	Number of targets accessed through the adapter.
Connected Targets (iSCSI)	Number of connected targets on an iSCSI adapter.
WWN (Fibre Channel)	World Wide Name formed according to Fibre Channel standards that uniquely identifies the FC adapter.
iSCSI Name (iSCSI)	Unique name formed according to iSCSI standards that identifies the iSCSI adapter.
iSCSI Alias (iSCSI)	A friendly name used instead of the iSCSI name.
IP Address (hardware iSCSI)	Address assigned to the iSCSI adapter.
Discovery Methods (iSCSI)	Discovery methods the iSCSI adapter uses to access iSCSI targets.
Devices	All storage devices or LUNs the adapter can access.
Paths	All paths the adapter uses to access storage devices.

View Storage Adapters Information

You can display storage adapters that your host uses and review their information.

Procedure

- 1 In Inventory, select **Hosts and Clusters**.
- 2 Select a host and click the **Configuration** tab.
- 3 In Hardware, select **Storage Adapters**.
- 4 To view details for a specific adapter, select the adapter from the Storage Adapters list.

- 5 To list all storage devices the adapter can access, click **Devices**.
- 6 To list all paths the adapter uses, click **Paths**.

Copy Storage Adapter Identifiers to the Clipboard

If your storage adapters use unique identifiers, such as an iSCSI Name or WWN, you can copy them to a clipboard directly from the UI.

Procedure

- 1 In Inventory, select **Hosts and Clusters**.
- 2 Select a host and click the **Configuration** tab.
- 3 In Hardware, select **Storage Adapters**.
- 4 Select the adapter from the Storage Adapters list.
- 5 In the Details panel, right-click the value in the name field, and select **Copy**.

Viewing Storage Devices

You can display all storage devices or LUNs available to the host, including all local and networked devices. If you use third-party multipathing plug-ins, the storage devices available through the plug-ins also appear on the list.

For each storage adapter, you can display a separate list of storage devices available just for this adapter.

Generally, when you review a list of storage devices, you see the following information.

Table 7-4. Storage Device Information

Storage Device Information	Description
Name	A friendly name that the ESX host assigns to the device based on the storage type and manufacturer. You can change this name to a name of your choice.
Identifier	A universally unique identifier that is intrinsic to the device.
Runtime Name	The name of the first path to the device.
LUN	The LUN number that shows the position of the LUN within the target.
Type	Type of device, for example, disk or CD-ROM.
Transport	Transportation protocol your host uses to access the device.
Capacity	Total capacity of the storage device.
Owner	The plug-in, such as the NMP or a third-party plug-in, that the host uses to manage the storage device.

Details for each storage device include the following:

- A path to the storage device in the `/vmfs/devices/` directory.
- Primary and logical partitions, including a VMFS datastore, if configured.

Display Storage Devices for a Host

You can display all storage devices or LUNs available to a host. If you use any third-party multipathing plugins, the storage devices available through the plugins also appear on the list.

Procedure

- 1 In Inventory, select **Hosts and Clusters**.
- 2 Select a host and click the **Configuration** tab.

- 3 In Hardware, select **Storage**.
- 4 Click **Devices**.
- 5 To view additional details about a specific device, select the device from the list.

Display Storage Devices for an Adapter

You can display a list of storage devices accessible to a specific storage adapter on the host.

Procedure

- 1 In Inventory, select **Hosts and Clusters**.
- 2 Select a host and click the **Configuration** tab.
- 3 In Hardware, select **Storage Adapters**.
- 4 Select the adapter from the Storage Adapters list.
- 5 Click **Devices**.

Copy Storage Device Identifiers to the Clipboard

A storage device identifier is a universally unique ID assigned to a storage device or LUN. Depending on the type of storage, different algorithms are used to create the identifier, and it can be long and complex. You can copy the storage device identifier directly from the UI.

Procedure

- 1 Display a list of storage devices.
- 2 Right-click a device and select **Copy identifier to clipboard**.

Displaying Datastores

You can display all datastores available to your hosts and analyze their properties.

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

- Created on an available storage device.
- Discovered when a host is added to the inventory. When you add a host to the inventory, the vSphere Client displays any datastores available to the host.

If your vSphere Client is connected to a vCenter Server system, you can see datastore information in the Datastores view. This view displays all datastores in the inventory, arranged by a datacenter. Through this view, you can organize datastores into folder hierarchies, create new datastores, edit their properties, or remove existing datastores.

This view is comprehensive and shows all information for your datastores including hosts and virtual machines using the datastores, storage reporting information, permissions, alarms, tasks and events, storage topology, and storage reports. Configuration details for each datastore on all hosts connected to this datastore are provided on the Configuration tab of the Datastores view.

NOTE The Datastores view is not available when the vSphere client connects directly to your host. In this case, review datastore information through the host storage configuration tab.

Generally, you can see the following datastore configuration details:

- Target storage device where the datastore is located
- File system that the datastore uses
- Location of the datastore

- Total capacity, including the used and available space
- Individual extents that the datastore spans and their capacity (VMFS datastores only)
- Paths used to access the storage device (VMFS datastores only)

Review Datastore Properties

You can display all datastores available to the hosts and analyze their properties.

Procedure

- 1 Display the host in the inventory.
- 2 Select a host in the inventory and click the **Configuration** tab.
- 3 In Hardware, select **Storage**.
- 4 Click the **Datastores** view.
- 5 To display details for a particular datastore, select the datastore from the list.

Configuring ESX Storage

The following topics contain information about configuring local SCSI storage devices, Fibre Channel SAN storage, iSCSI storage, and NFS storage.

This chapter includes the following topics:


- [“Local SCSI Storage,”](#) on page 83
- [“Fibre Channel Storage,”](#) on page 84
- [“iSCSI Storage,”](#) on page 84
- [“Storage Refresh and Rescan Operations,”](#) on page 94
- [“Create VMFS Datastores,”](#) on page 95
- [“Network Attached Storage,”](#) on page 95
- [“Creating a Diagnostic Partition,”](#) on page 97

Local SCSI Storage

Local storage uses a SCSI-based device such as your ESX host’s hard disk or any external dedicated storage system connected directly to your host.

[Figure 8-1](#) depicts a virtual machine using local SCSI storage.

Figure 8-1. Local Storage



In this example of a local storage topology, the ESX 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 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.


Fibre Channel Storage

ESX supports Fibre Channel adapters, which allow a host to connect to a SAN and see storage devices on the SAN.

You must install Fibre Channel (FC) adapters before the host can display FC storage devices.

Figure 8-2 depicts virtual machines using Fibre Channel storage.

Figure 8-2. Fibre Channel Storage



In this configuration, an ESX host 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 the host. You can access the LUNs and create a datastore for your storage needs. The datastore uses the VMFS format.

For specific information on setting up the FC SAN fabric and storage arrays to work with ESX, see the *Fibre Channel SAN Configuration Guide*.


iSCSI Storage

ESX supports iSCSI technology that allows your host 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 packets and transmitted to a remote device, or target, that stores the virtual disk.

To access remote targets, the host uses iSCSI initiators. Initiators transport SCSI requests and responses between the host and the target storage device on the IP network. ESX supports hardware-based and software-based iSCSI initiators.

You must configure iSCSI initiators for the host to access and display iSCSI storage devices.

Figure 8-3 depicts two virtual machines that use different types of iSCSI initiators.

Figure 8-3. iSCSI Storage

In the left example, the host uses the hardware iSCSI adapter to connect to the iSCSI storage system.

In the right example, the host is configured with the software iSCSI initiator. Using the software initiator, the host connects to the iSCSI storage through an existing network adapter.

iSCSI storage devices from the storage system become available to the host. You can access the storage devices and create VMFS datastores for your storage needs.

For specific information on setting up the iSCSI SAN fabric to work with ESX, see the *iSCSI SAN Configuration Guide*.

Setting Up Hardware iSCSI Initiators

With hardware-based iSCSI storage, you use a specialized third-party adapter capable of accessing iSCSI storage over TCP/IP. This iSCSI initiator handles all iSCSI and network processing and management for your ESX system.

You must install and configure the hardware iSCSI adapter for your host to be able to access the iSCSI storage device. For installation information, see vendor documentation.

View Hardware iSCSI Initiators

View an iSCSI hardware initiator to verify that it is correctly installed and ready for configuration.

Prerequisites

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

Procedure

- 1 Log in to the vSphere Client, and select a host 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 view.

The default details for the initiator appear, including the model, iSCSI name, iSCSI alias, IP address, and target and paths information.

- 4 Click **Properties**.

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

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

Change Name and IP Address for Hardware Initiators

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

Procedure

- 1 Log in to the vSphere Client, and select a host from the inventory panel.
- 2 Click the **Configuration** tab and click **Storage Adapters** in the Hardware panel.
- 3 Select the initiator to configure and click **Properties > Configure**.
- 4 To change the default iSCSI name for your initiator, enter the new name.

Make sure the name you enter is worldwide unique and properly formatted or some storage devices might not recognize the hardware iSCSI initiator.

- 5 (Optional) Enter the iSCSI alias.

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

- 6 Change the default IP settings.

You must change the default IP settings so that they are configured properly for the IP SAN. Work with your network administrator to determine the IP setting for the HBA.

- 7 Click **OK** to save your changes.

If you change the iSCSI name, it is used for new iSCSI sessions. For existing sessions, new settings are not used until logout and re-login.

Setting Up Software iSCSI Initiators

With the software-based iSCSI implementation, you can use standard network adapters to connect your ESX host to a remote iSCSI target on the IP network. The software iSCSI initiator that is built into ESX facilitates this connection by communicating with the network adapter through the network stack.

Before you configure the software iSCSI initiator, you must perform the following tasks:

- 1 Create a VMkernel port for physical network adapters.
- 2 Enable the software iSCSI initiator.
- 3 If you use multiple network adapters, activate multipathing on your host using the port binding technique.
For more information on port binding, see the *iSCSI SAN Configuration Guide*.
- 4 If needed, enable Jumbo Frames. Jumbo Frames must be enabled for each vSwitch through the vSphere CLI. Also, if you use an ESX host, you must create a VMkernel network interface enabled with Jumbo Frames.

See the *Networking* section for more information.

Networking Configuration for Software iSCSI Storage

Networking configuration for software iSCSI involves creating an iSCSI VMkernel port and mapping it to a physical NIC that handles iSCSI traffic.

Depending on the number of physical NICs you use for iSCSI traffic, the networking setup can be different:

- If you have one physical NIC, create one VMkernel port on a vSwitch and map the port to the NIC. VMware recommends that you designate a separate network adapter entirely for iSCSI. No additional network configuration steps are required.

For information on creating a port, see [“Create a VMkernel Port for Software iSCSI,”](#) on page 87.

- If you have two or more physical NICs for iSCSI, you can create multiple paths for the software iSCSI by using the port binding technique.

For more information on port binding, see the *iSCSI SAN Configuration Guide*.

Create a VMkernel Port for Software iSCSI

This procedure lets you connect the VMkernel, which runs services for iSCSI storage, to the physical network adapter.

Procedure

- 1 Log in to the vSphere Client and select the host from the inventory panel.
- 2 Click the **Configuration** tab and click **Networking**.
- 3 In the Virtual Switch view, click **Add Networking**.
- 4 Select **VMkernel** and click **Next**.
- 5 Select **Create a virtual switch** to create a new vSwitch.

If no adapters appear under **Create a virtual switch**, existing vSwitches are using all of the network adapters in the system. You can use an existing vSwitch for your iSCSI traffic.

- 6 Select an adapter you want to use for iSCSI traffic.

IMPORTANT Do not use iSCSI on 100Mbps or slower adapters.

- 7 Click **Next**.
- 8 Under **Port Group Properties**, enter a network label. Network label is a friendly name that identifies the VMkernel port that you are creating.
- 9 Click **Next**.
- 10 Specify the IP settings and click **Next**.
- 11 Review the information and click **Finish**.

What to do next

You can now enable your software initiator.

Enable the Software iSCSI Initiator

You must enable your software iSCSI initiator so that ESX can use it to access iSCSI storage.

Procedure

- 1 Log in to the vSphere 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 Select the iSCSI initiator to configure and click **Properties**.

- 4 Click **Configure**.

The **General Properties** dialog box displays the initiator's status, default name, and alias.

- 5 To enable the initiator, select **Enabled**.

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

Make sure the name you enter is worldwide unique and properly formatted or some storage devices might not recognize the software iSCSI initiator.

- 7 Click **OK** to save your changes.

If you change the iSCSI name, it is used for new iSCSI sessions. For existing sessions, new settings are not used until you logout and re-login.

Configuring Discovery Addresses for iSCSI Initiators

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

The ESX system supports these discovery methods:

Dynamic Discovery

Also known as Send Targets discovery. Each time the initiator contacts a specified iSCSI server, the initiator sends the Send Targets request to the server. The server responds by supplying a list of available targets to the initiator. The names and IP addresses of these targets appear on the **Static Discovery** tab. If you remove a static 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 host is rebooted.

Static Discovery

The initiator does not have to perform any discovery. The initiator has a list of targets it can contact and uses their IP addresses and target names to communicate with them.

Set Up Dynamic Discovery

With Dynamic Discovery, each time the initiator contacts a specified iSCSI server, it sends the Send Targets request to the server. The server responds by supplying a list of available targets to the initiator.

Procedure

- 1 Log in to the vSphere 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 Select the iSCSI initiator to configure and click **Properties**.

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

- 5 To add an address for the Send Targets discovery, click **Add**.

The **Add Send Targets Server** dialog box appears.

- 6 Enter the IP address or DNS name of the storage system and click **OK**.

After your host establishes the Send Targets session with this system, any newly discovered targets appear in the Static Discovery list.

- 7 To delete a specific Send Targets server, select it and click **Remove**.

After you remove a Send Targets server, it might still appear in the Inheritance field as the parent of static targets. This entry indicates where the static targets were discovered and does not affect the functionality.

NOTE You cannot change the IP address, DNS name, or port number of an existing Send Targets server. To make changes, delete the existing server and add a new one.

Set Up Static Discovery

With iSCSI initiators, in addition to the dynamic discovery method, you can use static discovery and manually enter information for the targets.

Procedure

- 1 Log in to the vSphere 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 Select the iSCSI initiator to configure and click **Properties**.
- 4 In the iSCSI Initiator Properties dialog box, click the **Static Discovery** tab.
The tab displays all dynamically discovered targets and any static targets already entered.
- 5 To add a target, click **Add** and enter the target's information.
- 6 To delete a specific target, select the target and click **Remove**.

NOTE You cannot change the IP address, DNS name, iSCSI target name, or port number of an existing target. To make changes, remove the existing target and add a new one.

Configuring CHAP Parameters for iSCSI Initiators

Because the IP networks that the iSCSI technology uses to connect to remote targets do not protect the data they transport, you must ensure security of the connection. One of the protocols that iSCSI implements is the Challenge Handshake Authentication Protocol (CHAP), which verifies the legitimacy of initiators that access targets on the network.

CHAP uses a three-way handshake algorithm to verify the identity of your host and, if applicable, of the iSCSI target when the host and target establish a connection. The verification is based on a predefined private value, or CHAP secret, that the initiator and target share.

ESX supports CHAP authentication at the adapter level. In this case, all targets receive the same CHAP name and secret from the iSCSI initiator. For software iSCSI, ESX also supports per-target CHAP authentication, which allows you to configure different credentials for each target to achieve greater level of security.

Choosing CHAP Authentication Method

ESX supports one-way CHAP for both hardware and software iSCSI, and mutual CHAP for software iSCSI only.

Before configuring CHAP, check whether CHAP is enabled at the iSCSI storage system and check the CHAP authentication method the system supports. If CHAP is enabled, enable it for your initiators, making sure that the CHAP authentication credentials match the credentials on the iSCSI storage.

ESX supports the following CHAP authentication methods:

One-way CHAP	In one-way, or unidirectional, CHAP authentication, the target authenticates the initiator, but the initiator does not authenticate the target.
Mutual CHAP (software iSCSI only)	In mutual, or bidirectional, CHAP authentication, an additional level of security enables the initiator to authenticate the target.

For software iSCSI only, you can set one-way CHAP and mutual CHAP for each initiator or at the target level. Hardware iSCSI supports CHAP only at the initiator level.

When you set the CHAP parameters, specify a security level for CHAP.

Table 8-1. CHAP Security Level

CHAP Security Level	Description	Supported
Do not use CHAP	The host does not use CHAP authentication. Select this option to disable authentication if it is currently enabled.	Software iSCSI Hardware iSCSI
Do not use CHAP unless required by target	The host prefers a non-CHAP connection, but can use a CHAP connection if required by the target.	Software iSCSI
Use CHAP unless prohibited by target	The host prefers CHAP, but can use non-CHAP connections if the target does not support CHAP.	Software iSCSI Hardware iSCSI
Use CHAP	The host requires successful CHAP authentication. The connection fails if CHAP negotiation fails.	Software iSCSI

Set Up CHAP Credentials for an iSCSI Initiator

For increased security, you can set up all targets to receive the same CHAP name and secret from the iSCSI initiator at the initiator level. By default, all discovery addresses or static targets inherit CHAP parameters that you set up at the initiator level.

Prerequisites

Before setting up CHAP parameters for software iSCSI, determine whether to configure one-way or mutual CHAP. Hardware iSCSI does not support mutual CHAP.

- In one-way CHAP, the target authenticates the initiator.
- In mutual CHAP, both the target and initiator authenticate each other. Make sure to use different secrets for CHAP and mutual CHAP.

When configuring CHAP parameters, make sure that they match the parameters on the storage side.

For software iSCSI, the CHAP name should not exceed 511 and the CHAP secret 255 alphanumeric characters. For hardware iSCSI, the CHAP name should not exceed 255 and the CHAP secret 100 alphanumeric characters.

Procedure

- 1 Log in to the vSphere 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 Select the iSCSI initiator to configure and click **Properties**.
- 4 On the **General** tab, click **CHAP**.

- 5 To configure one-way CHAP, under CHAP specify the following.
 - a Select one of the following options:
 - **Do not use CHAP unless required by target** (software iSCSI only)
 - **Use CHAP unless prohibited by target**
 - **Use CHAP** (software iSCSI only). To be able to configure mutual CHAP, you must select this option.
 - b Specify the CHAP name.
Make sure that the name you specify matches the name configured on the storage side.
 - To set the CHAP name to the iSCSI initiator name, select **Use initiator name**.
 - To set the CHAP name to anything other than the iSCSI initiator name, deselect **Use initiator name** and enter a name in the **Name** field.
 - c Enter a one-way CHAP secret to be used as part of authentication. Make sure to use the same secret that you enter on the storage side.
- 6 To configure mutual CHAP, first configure one-way CHAP by following directions in [Step 5](#).
Make sure to select **Use CHAP** as an option for one-way CHAP. Then, specify the following under **Mutual CHAP**:
 - a Select **Use CHAP**.
 - b Specify the mutual CHAP name.
 - c Enter the mutual CHAP secret. Make sure to use different secrets for the one-way CHAP and mutual CHAP.
- 7 Click **OK**.
- 8 Rescan the initiator.

If you change the CHAP or mutual CHAP parameters, they are used for new iSCSI sessions. For existing sessions, new settings are not used until you log out and login again.

Set Up CHAP Credentials for a Target

For software iSCSI, you can configure different CHAP credentials for each discovery address or static target.

When configuring CHAP parameters, make sure that they match the parameters on the storage side. For software iSCSI, the CHAP name should not exceed 511 and the CHAP secret 255 alphanumeric characters.

Prerequisites

Before setting up CHAP parameters for software iSCSI, determine whether to configure one-way or mutual CHAP.

- In one-way CHAP, the target authenticates the initiator.
- In mutual CHAP, both the target and initiator authenticate each other. Make sure to use different secrets for CHAP and mutual CHAP.

Procedure

- 1 Log in to the vSphere 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 Select the software iSCSI initiator to configure and click **Properties**.

- 4 Select either **Dynamic Discovery** tab or **Static Discovery** tab.
- 5 From the list of available targets, select a target you want to configure and click **Settings > CHAP**.
- 6 To configure one-way CHAP, under CHAP specify the following.
 - a Deselect **Inherit from parent**.
 - b Select one of the following options:
 - **Do not use CHAP unless required by target**
 - **Use CHAP unless prohibited by target**
 - **Use CHAP**. To be able to configure mutual CHAP, you must select this option.
 - c Specify the CHAP name.
Make sure that the name you specify matches the name configured on the storage side.
 - To set the CHAP name to the iSCSI initiator name, select **Use initiator name**.
 - To set the CHAP name to anything other than the iSCSI initiator name, deselect **Use initiator name** and enter a name in the **Name** field.
 - d Enter a one-way CHAP secret to be used as part of authentication. Make sure to use the same secret that you enter on the storage side.
- 7 To configure mutual CHAP, first configure one-way CHAP by following directions in [Step 6](#).
Make sure to select **Use CHAP** as an option for one-way CHAP. Then, specify the following under **Mutual CHAP**:
 - a Deselect **Inherit from parent**.
 - b Select **Use CHAP**.
 - c Specify the mutual CHAP name.
 - d Enter the mutual CHAP secret. Make sure to use different secrets for the one-way CHAP and mutual CHAP.
- 8 Click **OK**.
- 9 Rescan the initiator.

If you change the CHAP or mutual CHAP parameters, they are used for new iSCSI sessions. For existing sessions, new settings are not used until you log out and login again.

Disable CHAP

You can disable CHAP if your storage system does not require it.

If you disable CHAP on a system that requires CHAP authentication, existing iSCSI sessions remain active until you reboot your ESX host or the storage system forces a logout. After the session ends, you can no longer connect to targets that require CHAP.

Procedure

- 1 Open the CHAP Credentials dialog box.
- 2 For software iSCSI, to disable just the mutual CHAP, select **Do not use CHAP** under **Mutual CHAP**.
- 3 To disable one-way CHAP, select **Do not use CHAP** under **CHAP**.
The mutual CHAP, if set up, automatically turns to **Do not use CHAP** when you disable the one-way CHAP.
- 4 Click **OK**.

Configuring Additional Parameters for iSCSI

You might need to configure additional parameters for your iSCSI initiators. For example, some iSCSI storage systems require ARP (Address Resolution Protocol) redirection to move iSCSI traffic dynamically from one port to another. In this case, you must activate ARP redirection on your host.

Do not make any changes to the advanced iSCSI settings unless you are working with the VMware support team or otherwise have thorough information about the values to provide for the settings.

Table 8-2 lists advanced iSCSI parameters that you can configure using the vSphere Client. In addition, you can use the `vicfg-iscsi` vSphere CLI command to configure some of the advanced parameters. For information, see the *VMware vSphere Command-Line Interface Installation and Reference Guide*.

Table 8-2. Additional Parameters for iSCSI Initiators

Advanced Parameter	Description	Configurable On
Header Digest	Increases data integrity. When header digest is enabled, the system performs a checksum over each iSCSI Protocol Data Unit's (PDU's) header part and verifies using the CRC32C algorithm.	Software iSCSI
Data Digest	Increases data integrity. When data digest is enabled, the system performs a checksum over each PDU's data part and verifies using the CRC32C algorithm. NOTE Systems that use Intel Nehalem processors offload the iSCSI digest calculations for software iSCSI, thus reducing the impact on performance.	Software iSCSI
Maximum Outstanding R2T	Defines the R2T (Ready to Transfer) PDUs that can be in transition before an acknowledge PDU is received.	Software iSCSI
First Burst Length	Specifies the maximum amount of unsolicited data an iSCSI initiator can send to the target during the execution of a single SCSI command, in bytes.	Software iSCSI
Maximum Burst Length	Maximum SCSI data payload in a Data-In or a solicited Data-Out iSCSI sequence, in bytes.	Software iSCSI
Maximum Receive Data Segment Length	Maximum data segment length, in bytes, that can be received in an iSCSI PDU.	Software iSCSI
ARP Redirect	Allows storage systems to move iSCSI traffic dynamically from one port to another. ARP is required by storage systems that do array-based failover.	Hardware iSCSI (Configurable through vSphere CLI)
Delayed ACK	Allows systems to delay acknowledgment of received data packets.	Software iSCSI

Configure Advanced Parameters for iSCSI

The advanced iSCSI settings control such parameters as header and data digest, ARP redirection, delayed ACK, and so on. Generally, you do not need to change these settings because your ESX host works with the assigned predefined values.



CAUTION Do not make any changes to the advanced iSCSI settings unless you are working with the VMware support team or otherwise have thorough information about the values to provide for the settings.

Procedure

- 1 Log in to the vSphere Client, and select a host from the inventory panel.
- 2 Click **Configuration** tab and click **Storage Adapters**.
- 3 Select the iSCSI initiator to configure and click **Properties**.

- 4 To configure advanced parameters at the initiator level, on the **General** tab, click **Advanced**. Proceed to [Step 6](#).
- 5 Configure advanced parameters at the target level.
At the target level, advanced parameters can be configured only for software iSCSI.
 - a Select either the **Dynamic Discovery** tab or **Static Discovery** tab.
 - b From the list of available targets, select a target to configure and click **Settings > Advanced**.
- 6 Enter any required values for the advanced parameters you want to modify and click **OK** to save your changes.

Storage Refresh and Rescan Operations

The refresh operation updates the datastore lists and storage information, such as the datastore capacity, displayed in the vSphere Client. When you make changes in your ESX host or SAN configuration, you need to use the rescan operation.

You can rescan all adapters on your host. If the changes you make are isolated to a specific adapter, rescan only this adapter. If your vSphere Client is connected to a vCenter Server system, you can rescan adapters on all hosts managed by the vCenter Server system.

Perform a rescan each time you make one of the following changes.

- Create new LUNs on a SAN.
- Change the path masking on a host.
- Reconnect a cable.
- Make a change to a host in a cluster.

IMPORTANT Do not rescan when a path is unavailable. If one path fails, another takes over and your system continues to be fully functional. If, however, you rescan at a time when a path is not available, the host removes the path from its list of paths to the device. The path cannot be used by the host until the next time a rescan is performed while the path is active.

Rescan Storage Adapters

When you make changes in your ESX host or SAN configuration, you might need to rescan your storage adapters. You can rescan all adapters on your host. If the changes you make are isolated to a specific adapter, rescan only this adapter.

Use this procedure if you want to limit the rescan to a particular host or an adapter on the host. If you want to rescan adapters on all hosts managed by your vCenter Server system, you can do so by right-clicking a datacenter, cluster, or folder that contains the hosts and selecting **Rescan for Datastores**.

Procedure

- 1 In the vSphere Client, select a host and click the **Configuration** tab.
- 2 In the Hardware panel, select **Storage Adapters**, and click **Rescan** above the Storage Adapters panel.
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 device 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.

Create VMFS Datastores

VMFS datastores serve as repositories for virtual machines. You can set up VMFS datastores on any SCSI-based storage devices that the host discovers.

Prerequisites

Before creating datastores, you must install and configure any adapters that your storage requires. Rescan the adapters to discover newly added storage devices.

Procedure

- 1 Log in to the vSphere Client and select the host from the Inventory panel.
- 2 Click the **Configuration** tab and click **Storage** in the Hardware panel.
- 3 Click **Datastores** and click **Add Storage**.
- 4 Select the **Disk/LUN** storage type and click **Next**.
- 5 Select a device to use for your datastore and click **Next**.

NOTE Select the device that does not have a datastore name displayed in the VMFS Label column. If a name is present, the device contains a copy of an existing VMFS datastore.

If the disk you are formatting is blank, the Current Disk Layout page automatically presents the entire disk space 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.

Option	Description
Use all available partitions	Dedicates the entire disk or LUN to a single VMFS datastore. If you select this option, all file systems and data currently stored on this device is destroyed.
Use free space	Deploys a VMFS datastore in the remaining free space of the disk.

- 7 Click **Next**.
- 8 In the Properties page, enter a datastore name and click **Next**.
- 9 If needed, adjust the file system and capacity values.
By default, the entire free space on the storage device is available.
- 10 Click **Next**.
- 11 In the Ready to Complete page, review the datastore configuration information and click **Finish**.

A datastore on the SCSI-based storage device is created. If you use the vCenter Server system to manage your hosts, the newly created datastore is automatically added to all hosts.

Network Attached Storage

ESX supports using NAS through the NFS protocol. The NFS protocol enables communication between an NFS client and an NFS server.


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

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

NOTE ESX does not support the delegate user functionality that enables access to NFS volumes using non-root credentials.

Figure 8-4 depicts a virtual machine using the NFS volume to store its files. In this configuration, the host connects to the NFS server, which stores the virtual disk files, through a regular network adapter.

Figure 8-4. NFS Storage



The virtual disks that you create on NFS-based datastores use a disk format dictated by the NFS server, typically a thin format that requires on-demand space allocation. If the virtual machine runs out of space while writing to this disk, the vSphere 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.



CAUTION When your host accesses a virtual machine disk file on an NFS-based datastore, a `.lck-XXX` lock file is generated in the same directory where the disk file resides to prevent other 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 Datastores as Repositories for Commonly Used Files

In addition to storing virtual disks on NFS datastores, you can also use NFS as a central repository for ISO images, virtual machine templates, and so on.

To use NFS as a shared repository, you create a directory on the NFS server and then mount it as a datastore on all hosts. If you use the datastore for ISO images, you can connect the virtual machine's CD-ROM device to an ISO file on the datastore and install a guest operating system from the ISO file.

For information on configuring virtual machines, see *Basic System Administration*.

NOTE If the underlying NFS volume, on which the files are stored, is read-only, make sure that the volume is exported as a read-only share by the NFS server, or configure it as a read-only datastore on the ESX host. Otherwise, the host considers the datastore to be read-write and might not be able to open the files.

Create an NFS-Based Datastore

You can use the Add Storage wizard to mount an NFS volume and use it as if it were a VMFS datastore.

Prerequisites

Because NFS requires network connectivity to access data stored on remote servers, before configuring NFS, you must first configure VMkernel networking.

Procedure

- 1 Log in to the vSphere Client and select the host from the Inventory panel.
- 2 Click the **Configuration** tab and click **Storage** in the Hardware panel.
- 3 Click **Datastores** and 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.

NOTE When you mount the same NFS volume on different hosts, make sure that the server and folder names are identical across the hosts. If the names do not match exactly, for example, if you enter **share** as the folder name on one host and **/share** on the other, the hosts see the same NFS volume as two different datastores. This might result in a failure of such features as vMotion.

- 6 (Optional) Select **Mount NFS read only** if the volume is exported as read only by the NFS server.
- 7 Click **Next**.
- 8 In the Network File System Summary page, review the configuration options and click **Finish**.

Creating a Diagnostic Partition

To run successfully, your host must 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 iSCSI initiator.

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



CAUTION If two hosts that share a diagnostic partition fail and save core dumps to the same slot, the core dumps might be lost. To collect core dump data, reboot a host and extract log files immediately after the host fails. However, if another host fails before you collect the diagnostic data of the first host, the second host will fail to save the core dump.

With the ESX host, you typically create a diagnostic partition when installing ESX by selecting **Recommended Partitioning**. The installer automatically creates a diagnostic partition for your host. If you select **Advanced Partitioning** and choose not to specify the diagnostic partition during installation, you can configure it using the Add Storage wizard.

Create a Diagnostic Partition

You can create a diagnostic partition on your host.

Procedure

- 1 Log in to the vSphere Client and select the host from the Inventory panel.
- 2 Click the **Configuration** tab and click **Storage** in the Hardware panel.
- 3 Click **Datastores** and click **Add Storage**.
- 4 Select **Diagnostic** and click **Next**.

If you do not see **Diagnostic** as an option, the host already has a diagnostic partition.

You can query and scan the host's diagnostic partition using the `vicfg-dumppart -l` command on the vSphere CLI.

- 5 Specify the type of diagnostic partition.

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

- 6 Click **Next**.
- 7 Select the device to use for the diagnostic partition and click **Next**.
- 8 Review the partition configuration information and click **Finish**.

Managing Storage

After you create datastores, you can change their properties, use folders to group datastores based on your business needs, or delete unused datastores. You might also need to set up multipathing for your storage or resignature datastore copies.

This chapter includes the following topics:

- [“Managing Datastores,”](#) on page 99
- [“Changing VMFS Datastore Properties,”](#) on page 101
- [“Managing Duplicate VMFS Datastores,”](#) on page 103
- [“Using Multipathing with ESX,”](#) on page 105
- [“Thin Provisioning,”](#) on page 113
- [“Turn Off vCenter Server Storage Filters,”](#) on page 116

Managing Datastores

An ESX system uses datastores to store all files associated with its virtual machines. After you create datastores, you can manage them by performing a number of tasks.

A datastore is a logical storage unit that 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.

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

- Discovered when a host is added to the inventory. The vSphere Client displays any datastores that the host can recognize.
- Created on an available storage device using the **Add Storage** command.

After the datastores are created, you can use them to store virtual machine files. You can manage them by renaming, removing, and setting access control permissions. In addition, you can group datastores to organize them and set the same permissions across the group at one time.

For information on setting access control permissions on a datastore, see the *vSphere Client Help*.

Rename Datastores

You can change the name of an existing datastore.

Procedure

- 1 Display the datastores.

- 2 Right-click the datastore to rename and select **Rename**.
- 3 Type a new datastore name.

If you use the vCenter Server system to manage your hosts, the new name appears on all hosts that have access to the datastore.

Group Datastores

If you use the vCenter Server system to manage your hosts, group datastores into folders. This allows you to organize your datastores according to business practices and to assign the same permissions and alarms on the datastores in the group at one time.

Procedure

- 1 Log in to the vSphere Client.
- 2 If necessary, create the datastores.
For details, see the vSphere Client Help.
- 3 In the Inventory panel, choose **Datastores**.
- 4 Select the datacenter containing the datastores to group.
- 5 In the shortcut menu, click the **New Folder** icon.
- 6 Give the folder a descriptive name.
- 7 Click and drag each datastore onto the folder.

Delete Datastores

You can delete any type of VMFS datastore, including copies that you have mounted without resignaturing. When you delete a datastore, it is destroyed and disappears from all hosts that have access to the datastore.

Prerequisites

Before deleting a datastore, remove all virtual machines from the datastore. Make sure that no other host is accessing the datastore.

Procedure

- 1 Display the datastores.
- 2 Right-click the datastore to delete and click **Delete**.
- 3 Confirm that you want to delete the datastore.

Unmount Datastores

When you unmount a datastore, it remains intact, but can no longer be seen from the hosts that you specify. It continues to appear on other hosts, where it remains mounted.

You can unmount only the following types of datastores:

- NFS datastores
- VMFS datastore copies mounted without resignaturing

Procedure

- 1 Display the datastores.
- 2 Right-click the datastore to unmount and select **Unmount**.

- 3 If the datastore is shared, specify which hosts should no longer access the datastore.
 - a If needed, deselect the hosts where you want to keep the datastore mounted.
By default, all hosts are selected.
 - b Click **Next**.
 - c Review the list of hosts from which to unmount the datastore, and click **Finish**.
- 4 Confirm that you want to unmount the datastore.

Changing VMFS Datastore Properties

After you create a VMFS-based datastore, you can modify it. For example, you can increase it if you need more space. If you have VMFS-2 datastores, you can upgrade them to VMFS-3 format.

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

You cannot reformat a VMFS datastore that a remote host is using. If you attempt to, a warning appears that specifies the name of the datastore in use and the host that is using it. This warning also appears in the VMkernel and vmkwarning log files.

Depending on whether your vSphere Client is connected to a vCenter Server system or directly to a host, different ways to access the Datastore Properties dialog box exist.

- vCenter Server only. To access the Datastore Properties dialog box, select the datastore from the inventory, click the **Configuration** tab, and click **Properties**.
- vCenter Server and ESX/ESXi host. To access the Datastore Properties dialog box, select a host from the inventory, click the **Configuration** tab and click **Storage**. From the Datastores view, select the datastore to modify and click **Properties**.

Increase VMFS Datastores

When you need to create new virtual machines on a datastore, or when the virtual machines running on this datastore require more space, you can dynamically increase the capacity of a VMFS datastore.

Use one of the following methods:

- Add a new extent. An extent is a partition on a LUN. You can add a new extent to any existing VMFS datastore. The datastore can stretch over multiple extents, up to 32.

NOTE You cannot add a local extent to a datastore located on a SAN LUN.

- Grow an extent in an existing VMFS datastore. Only extents with free space immediately after them are expandable. As a result, rather than adding the new extent, you can grow the existing extent so that it fills the available adjacent capacity.

NOTE If a shared datastore has powered on virtual machines and becomes 100% full, you can increase the datastore's capacity only from the host, with which the powered on virtual machines are registered.

Procedure

- 1 Log in to the vSphere Client and select a host from the Inventory panel.
- 2 Click the **Configuration** tab and click **Storage**.
- 3 From the Datastores view, select the datastore to increase and click **Properties**.
- 4 Click **Increase**.

- 5 Select a device from the list of storage devices and click **Next**.
 - If you want to add a new extent, select the device for which the Expandable column reads No.
 - If you want to expand an existing extent, select the device for which the Expandable column reads Yes.

- 6 Select a configuration option from the bottom panel.

Depending on the current layout of the disk and on your previous selections, the options you see might vary.

Option	Description
Use free space to add new extent	Adds the free space on this disk as a new datastore extent.
Use free space to expand existing extent	Grows an existing extent to a required capacity.
Use free space	Deploys an extent in the remaining free space of the disk. This option is available only when adding an extent.
Use all available partitions	Dedicates the entire disk to a single datastore extent. This option is available only when adding an extent and when the disk you are formatting is not blank. The disk is reformatted, and the datastores and any data that it contains are erased.

- 7 Set the capacity for the extent.

By default, the entire free space on the storage device is available.
- 8 Click **Next**.
- 9 Review the proposed layout and the new configuration of your datastore, and click **Finish**.

What to do next

After you grow an extent in a shared VMFS datastore, refresh the datastore on each host that can access this datastore, so that the vSphere Client can display the correct datastore capacity for all hosts.

Upgrade Datastores

ESX 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 write to them. To have complete access to the files, upgrade VMFS-2 to VMFS-3.

When you upgrade VMFS-2 to VMFS-3, the ESX file-locking mechanism ensures that no remote host or local process is accessing the VMFS datastore being converted. Your host 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 that you plan to upgrade.
- Back up the VMFS-2 volume.
- Be sure that no powered on virtual machines are using the VMFS-2 volume.
- Be sure that no other ESX host is accessing the VMFS-2 volume.

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 upgrade the VMFS-2 file system, its file block size must not exceed 8MB.

Procedure

- 1 Log in to the vSphere Client and select a host 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.

Managing Duplicate VMFS Datastores

When a LUN contains a VMFS datastore copy, you can mount the datastore with the existing signature or assign a new signature.

Each VMFS datastore created in a LUN has a unique UUID that is stored in the file system superblock. When the LUN is replicated or snapshotted, the resulting LUN copy is identical, byte-for-byte, with the original LUN. As a result, if the original LUN contains a VMFS datastore with UUID X, the LUN copy appears to contain an identical VMFS datastore, or a VMFS datastore copy, with exactly the same UUID X.

ESX can determine whether a LUN contains the VMFS datastore copy, and either mount the datastore copy with its original UUID or change the UUID, thus resignaturing the datastore.

Mounting VMFS Datastores with Existing Signatures

You might not have to resignature a VMFS datastore copy. You can mount a VMFS datastore copy without changing its signature.

For example, you can maintain synchronized copies of virtual machines at a secondary site as part of a disaster recovery plan. In the event of a disaster at the primary site, you can mount the datastore copy and power on the virtual machines at the secondary site.

IMPORTANT You can mount a VMFS datastore only if it does not collide with an already mounted VMFS datastore that has the same UUID.

When you mount the VMFS datastore, ESX allows both reads and writes to the datastore residing on the LUN copy. The LUN copy must be writable. The datastore mounts are persistent and valid across system reboots.

Because ESX does not allow you to resignature the mounted datastore, unmount the datastore before resignaturing.

Mount a VMFS Datastore with an Existing Signature

If you do not need to resignature a VMFS datastore copy, you can mount it without changing its signature.

Prerequisites

Before you mount a VMFS datastore, perform a storage rescan on your host so that it updates its view of LUNs presented to it.

Procedure

- 1 Log in to the vSphere 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 From the list of LUNs, select the LUN that has a datastore name displayed in the VMFS Label column and click **Next**.

The name present in the VMFS Label column indicates that the LUN is a copy that contains a copy of an existing VMFS datastore.

- 6 Under Mount Options, select **Keep Existing Signature**.

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

What to do next

If you later want to resignature the mounted datastore, you must unmount it first.

Resignaturing VMFS Copies

Use datastore resignaturing to retain the data stored on the VMFS datastore copy. When resignaturing a VMFS copy, ESX assigns a new UUID and a new label to the copy, and mounts the copy as a datastore distinct from the original.

The default format of the new label assigned to the datastore is `snap-<snapID>-<oldLabel>`, where *<snapID>* is an integer and *<oldLabel>* is the label of the original datastore.

When you perform datastore resignaturing, consider the following points:

- Datastore resignaturing is irreversible.
- The LUN copy that contains the VMFS datastore that you resignature is no longer treated as a LUN copy.
- A spanned datastore can be resignatured only if all its extents are online.
- The resignaturing process is crash and fault tolerant. If the process is interrupted, you can resume it later.
- You can mount the new VMFS datastore without a risk of its UUID colliding with UUIDs of any other datastore, such as an ancestor or child in a hierarchy of LUN snapshots.

Resignature a VMFS Datastore Copy

Use datastore resignaturing if you want to retain the data stored on the VMFS datastore copy.

Prerequisites

To resignature a mounted datastore copy, first unmount it.

Before you resignature a VMFS datastore, perform a storage rescan on your host so that the host updates its view of LUNs presented to it and discovers any LUN copies.

Procedure

- 1 Log in to the vSphere 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 From the list of LUNs, select the LUN that has a datastore name displayed in the VMFS Label column and click **Next**.

The name present in the VMFS Label column indicates that the LUN is a copy that contains a copy of an existing VMFS datastore.

- 6 Under Mount Options, select **Assign a New Signature** and click **Next**.
- 7 In the Ready to Complete page, review the datastore configuration information and click **Finish**.

What to do next

After resignaturing, you might have to do the following:

- If the resignatured datastore contains virtual machines, update references to the original VMFS datastore in the virtual machine files, including `.vmx`, `.vmdk`, `.vmsd`, and `.vmsn`.
- To power on virtual machines, register them with vCenter Server.

Using Multipathing with ESX

To maintain a constant connection between an ESX host and its storage, ESX supports multipathing. Multipathing is a technique that lets you use more than one physical path for transferring data between the ESX host and the external storage device.

In case of a failure of an element in the SAN network, such as an HBA, switch, or cable, ESX can fail over to another physical path. In addition to path failover, multipathing offers load balancing, which redistributes I/O loads between multiple paths, thus reducing or removing potential bottlenecks.

Managing Multiple Paths

To manage storage multipathing, ESX uses a special VMkernel layer, Pluggable Storage Architecture (PSA). The PSA is an open modular framework that coordinates the simultaneous operation of multiple multipathing plugins (MPPs).

The VMkernel multipathing plugin that ESX provides by default is the VMware Native Multipathing Plugin (NMP). The NMP is an extensible module that manages subplugins. There are two types of NMP subplugins, Storage Array Type Plugins (SATPs), and Path Selection Plugins (PSPs). SATPs and PSPs can be built-in and provided by VMware, or can be provided by a third party.


If more multipathing functionality is required, a third party can also provide an MPP to run in addition to, or as a replacement for, the default NMP.

When coordinating the VMware NMP and any installed third-party MPPs, the PSA performs the following tasks:

- Loads and unloads multipathing plugins.
- Hides virtual machine specifics from a particular plugin.
- Routes I/O requests for a specific logical device to the MPP managing that device.
- Handles I/O queuing to the logical devices.
- Implements logical device bandwidth sharing between virtual machines.
- Handles I/O queueing to the physical storage HBAs.
- Handles physical path discovery and removal.
- Provides logical device and physical path I/O statistics.

As [Figure 9-1](#) illustrates, multiple third-party MPPs can run in parallel with the VMware NMP. The third-party MPPs can replace the behavior of the NMP and take complete control of the path failover and the load-balancing operations for specified storage devices.

Figure 9-1. Pluggable Storage Architecture



The multipathing modules perform the following operations:

- Manage physical path claiming and unclaiming.
- Manage creation, registration, and deregistration of logical devices.
- Associate physical paths with logical devices.
- Process I/O requests to logical devices:
 - Select an optimal physical path for the request.
 - Depending on a storage device, perform specific actions necessary to handle path failures and I/O command retries.
- Support management tasks, such as abort or reset of logical devices.

VMware Multipathing Module

By default, ESX provides an extensible multipathing module called the Native Multipathing Plugin (NMP).

Generally, the VMware NMP supports all storage arrays listed on the VMware storage HCL and provides a default path selection algorithm based on the array type. The NMP associates a set of physical paths with a specific storage device, or LUN. The specific details of handling path failover for a given storage array are delegated to a Storage Array Type Plugin (SATP). The specific details for determining which physical path is used to issue an I/O request to a storage device are handled by a Path Selection Plugin (PSP). SATPs and PSPs are sub-plugins within the NMP module.

VMware SATPs

Storage Array Type Plugins (SATPs) run in conjunction with the VMware NMP and are responsible for array-specific operations.

ESX offers an SATP for every type of array that VMware supports. These SATPs include an active/active SATP and active/passive SATP for non-specified storage arrays, and the local SATP for direct-attached storage. Each SATP accommodates special characteristics of a certain class of storage arrays and can perform the array-specific operations required to detect path state and to activate an inactive path. As a result, the NMP module can work with multiple storage arrays without having to be aware of the storage device specifics.

After the NMP determines which SATP to call for a specific storage device and associates the SATP with the physical paths for that storage device, the SATP implements the tasks that include the following:

- Monitors health of each physical path.
- Reports changes in the state of each physical path.
- Performs array-specific actions necessary for storage fail-over. For example, for active/passive devices, it can activate passive paths.

VMware PSPs

Path Selection Plugins (PSPs) run in conjunction with the VMware NMP and are responsible for choosing a physical path for I/O requests.

The VMware NMP assigns a default PSP for every logical device based on the SATP associated with the physical paths for that device. You can override the default PSP.

By default, the VMware NMP supports the following PSPs:

Most Recently Used (MRU)

Selects the path the ESX host used most recently to access the given device. If this path becomes unavailable, the host switches to an alternative path and continues to use the new path while it is available.

Fixed

Uses the designated preferred path, if it has been configured. Otherwise, it uses the first working path discovered at system boot time. If the host cannot use the preferred path, it selects a random alternative available path. The host automatically reverts back to the preferred path as soon as that path becomes available.

NOTE With active-passive arrays that have a **Fixed** path policy, path thrashing might be a problem.

Round Robin (RR)

Uses a path selection algorithm that rotates through all available paths enabling load balancing across the paths.

VMware NMP Flow of I/O

When a virtual machine issues an I/O request to a storage device managed by the NMP, the following process takes place.


- 1 The NMP calls the PSP assigned to this storage device.
- 2 The PSP selects an appropriate physical path on which to issue the I/O.
- 3 If the I/O operation is successful, the NMP reports its completion.
- 4 If the I/O operation reports an error, the NMP calls an appropriate SATP.
- 5 The SATP interprets the I/O command errors and, when appropriate, activates inactive paths.
- 6 The PSP is called to select a new path on which to issue the I/O.

Multipathing with Local Storage and Fibre Channel SANs

In a simple multipathing local storage topology, you can use one ESX host, which has two HBAs. The ESX host connects to a dual-port local storage system through two cables. This configuration ensures fault tolerance if one of the connection elements between the ESX host and the local storage system fails.

To support path switching with FC SAN, the ESX host typically has two or more HBAs available from which the storage array can be reached using one or more switches. Alternatively, the setup can include one HBA and two storage processors so that the HBA can use a different path to reach the disk array.

In [Figure 9-2](#), 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 9-2. 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. ESX supports HBA and SP failover with its multipathing capability.

Multipathing with iSCSI SAN


With iSCSI storage, you can take advantage of the multipathing support that the IP network offers. In addition, ESX supports host-based multipathing for both hardware and software iSCSI initiators.

ESX can use multipathing support built into the IP network, which allows the network to perform routing. 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.

ESX also supports host-based multipathing.

With the hardware iSCSI, the host can have two or more hardware iSCSI adapters and use them as different paths to reach the storage system.


As [Figure 9-3](#) illustrates, the host has two hardware iSCSI adapters, HBA1 and HBA2, that provide two physical paths to the storage system. Multipathing plugins on the host, whether the VMkernel NMP or any third-party MPPs, have access to the paths by default and can monitor the health of each physical path. If, for example, HBA1 or the link between HBA1 and the network fails, the multipathing plugins can switch the path over to HBA2.

Figure 9-3. Hardware iSCSI and Failover

With the software iSCSI, as [Figure 9-4](#) shows, you can use multiple NICs that provide failover and load-balancing capabilities for iSCSI connections between the host and storage systems.

For this setup, because multipathing plugins do not have direct access to the physical NICs on your host, you must first connect each physical NIC to a separate VMkernel port. You then associate all VMkernel ports with the software iSCSI initiator using a port binding technique. As a result, each VMkernel port connected to a separate NIC becomes a different path that the iSCSI storage stack and its storage-aware multipathing plugins can use.

For more information on this setup, see the *iSCSI SAN Configuration Guide*.

Figure 9-4. Software iSCSI and Failover

Path Scanning and Claiming

When you start your ESX host or rescan your storage adapter, the host discovers all physical paths to storage devices available to the host. Based on a set of claim rules defined in the `/etc/vmware/esx.conf` file, the host determines which multipathing plugin (MPP) should claim the paths to a particular device and become responsible for managing the multipathing support for the device.

By default, the host performs a periodic path evaluation every 5 minutes causing any unclaimed paths to be claimed by the appropriate MPP.

The claim rules are numbered. For each physical path, the host runs through the claim rules starting with the lowest number first. The attributes of the physical path are compared to the path specification in the claim rule. If there is a match, the host assigns the MPP specified in the claim rule to manage the physical path. This continues until all physical paths are claimed by corresponding MPPs, either third-party multipathing plugins or the native multipathing plugin (NMP).

For the paths managed by the NMP module, a second set of claim rules is applied. These rules determine which SATP should be used to manage the paths from a specific array type, and which PSP is to be used for each storage device. For example, for a storage device that belongs to the EMC CLARiiON CX storage family, the default SATP is `VMW_SATP_CX` and the default PSP is Most Recently Used.

Use the vSphere Client to view which SATP and PSP the host is using for a specific storage device and the status of all available paths for this storage device. If needed, you can change the default VMware PSP using the vSphere Client. To change the default SATP, you need to modify claim rules using the vSphere CLI.

For detailed descriptions of the commands available to manage PSA, see the *vSphere Command-Line Interface Installation and Reference Guide*.

Viewing the Paths Information

Use the vSphere Client to determine which SATP and PSP the ESX host uses for a specific storage device and the status of all available paths for this storage device. You can access the path information from both, the Datastores and Devices views. For datastores, you review the paths that connect to the device the datastore is deployed on.

The path information includes the SATP assigned to manage the device, the path selection policy (PSP), and a list of paths with their physical characteristics, such as an adapter and target each path uses, and the status of each path. The following path status information can appear:

Active	Paths available for issuing I/O to a LUN. A single or multiple working paths currently used for transferring data are marked as Active (I/O).
NOTE For hosts that run ESX 3.5 or earlier, the term active means the only path that the host is using to issue I/O to a LUN.	
Standby	The path is operational and can be used for I/O if active paths fail.
Disabled	The path is disabled and no data can be transferred.
Broken	The software cannot connect to the disk through this path.

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 Preferred column.

View Datastore Paths

Use the vSphere Client to review the paths that connect to storage devices the datastores are deployed on.

Procedure

- 1 Log in to the vSphere Client and select a server from the inventory panel.
- 2 Click the **Configuration** tab and click **Storage** in the Hardware panel.
- 3 Click **Datastores** under View.
- 4 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 device and whether any of them are broken or disabled.
- 5 Click **Properties > Manage Paths** to open the Manage Paths dialog box.
You can use the Manage Paths dialog box to enable or disable your paths, set multipathing policy, and specify the preferred path.

View Storage Device Paths

Use the vSphere Client to view which SATP and PSP the host uses for a specific storage device and the status of all available paths for this storage device.

Procedure

- 1 Log in to the vSphere Client and select a server from the inventory panel.
- 2 Click the **Configuration** tab and click **Storage** in the Hardware panel.
- 3 Click **Devices** under View.
- 4 Click **Manage Paths** to open the Manage Paths dialog box.

Setting a Path Selection Policy

For each storage device, the ESX host sets the path selection policy based on the claim rules defined in the `/etc/vmware/esx.conf` file.

By default, VMware supports the following path selection policies. If you have a third-party PSP installed on your host, its policy also appears on the list.

Fixed (VMware)	The host always uses the preferred path to the disk when that path is available. If the host cannot access the disk through the preferred path, it tries the alternative paths. The default policy for active-active storage devices is Fixed.
Most Recently Used (VMware)	The host uses a path to the disk until the path becomes unavailable. When the path becomes unavailable, the host selects one of the alternative paths. The host does not revert back to the original path when that path becomes available again. There is no preferred path setting with the MRU policy. MRU is the default policy for active-passive storage devices and is required for those devices.
Round Robin (VMware)	The host uses an automatic path selection algorithm rotating through all available paths. This implements load balancing across all the available physical paths. Load balancing is the process of spreading server I/O requests across all available host paths. The goal is to optimize performance in terms of throughput (I/O per second, megabytes per second, or response times).

[Table 9-1](#) summarizes how the behavior of host changes, depending on the type of array and the failover policy.

Table 9-1. Path Policy Effects

Policy/Controller	Active-Active	Active-Passive
Most Recently Used	Administrator action is required to fail back after path failure.	Administrator action is required to fail back after path failure.
Fixed	VMkernel resumes using the preferred path when connectivity is restored.	VMkernel attempts to resume using the preferred path. This can cause path thrashing or failure when another SP now owns the LUN.
Round Robin	No fail back.	Next path in round robin scheduling is selected.

Change the Path Selection Policy

Generally, you do not have to change the default multipathing settings your host uses for a specific storage device. However, if you want to make any changes, you can use the Manage Paths dialog box to modify a path selection policy and specify the preferred path for the Fixed policy.

Procedure

- 1 Open the Manage Paths dialog box either from the Datastores or Devices view.

- 2 Select a path selection policy.

By default, VMware supports the following path selection policies. If you have a third-party PSP installed on your host, its policy also appears on the list.

- **Fixed (VMware)**
- **Most Recently Used (VMware)**
- **Round Robin (VMware)**

- 3 For the fixed policy, specify the preferred path by right-clicking the path you want to assign as the preferred path, and selecting **Preferred**.
- 4 Click **OK** to save your settings and exit the dialog box.

Disable Paths

You can temporarily disable paths for maintenance or other reasons. You can do so using the vSphere Client.

Procedure

- 1 Open the Manage Paths dialog box either from the Datastores or Devices view.
- 2 In the Paths panel, right-click the path to disable, and select **Disable**.
- 3 Click **OK** to save your settings and exit the dialog box.

You can also disable a path from the adapter's Paths view by right-clicking the path in the list and selecting **Disable**.

Thin Provisioning

When you create a virtual machine, a certain amount of storage space on a datastore is provisioned or allocated to the virtual disk files.

By default, ESX offers a traditional storage provisioning method during creation in which you estimate how much storage the virtual machine will need for its entire life cycle, provision a fixed amount of storage space to its virtual disk, and have the entire provisioned space committed to the virtual disk. A virtual disk that immediately occupies the entire provisioned space is called a thick disk. Creating virtual disks in thick format can lead to underutilization of datastore capacity, because large amounts of storage space, pre-allocated to individual virtual machines, might remain unused.

To help avoid over-allocating storage space and save storage, ESX supports thin provisioning, which lets you, in the beginning, use just as much storage capacity as currently needed and then add the required amount of storage space at a later time. Using the ESX thin provisioning feature, you can create virtual disks in a thin format. For a thin virtual disk, ESX provisions the entire space required for the disk's current and future activities, but commits only as much storage space as the disk needs for its initial operations.

About Virtual Disk Formats

When you perform certain virtual machine management operations, such as create a virtual disk, clone a virtual machine to a template, or migrate a virtual machine, you can specify a format for the virtual disk file.

The following disk formats are supported. You cannot specify the disk format if the disk resides on an NFS datastore. The NFS server determines the allocation policy for the disk.

Thin Provisioned Format Use this format to save storage space. For the thin disk, you provision as much datastore space as the disk would require based on the value you enter for the disk size. However, the thin disk starts small and at first, uses only as much datastore space as the disk actually needs for its initial operations.

NOTE If a virtual disk supports clustering solutions such as Fault Tolerance, you cannot make the disk thin.

If the thin disk needs more space later, it can grow to its maximum capacity and occupy the entire datastore space provisioned to it. Also, you can manually convert the thin disk into thick.

Thick Format This is the default virtual disk format. The thick virtual disk does not change its size and from the very beginning occupies the entire datastore space provisioned to it. Thick format does not zero out the blocks in the allocated space. It is not possible to convert the thick disk into thin.

Create Thin Provisioned Virtual Disks

When you need to save storage space, you can create a virtual disk in thin provisioned format. The thin provisioned virtual disk starts small and grows as more disk space is required.

This procedure assumes that you are creating a typical or custom virtual machine using the New Virtual Machine wizard.

Prerequisites

You can create thin disks only on the datastores that support thin provisioning. If a disk resides on an NFS datastore, you cannot specify the disk format because the NFS server determines the allocation policy for the disk.

Procedure

- ◆ In the Create a Disk dialog box, select **Allocate and commit space on demand (Thin Provisioning)**.

A virtual disk in thin format is created. If you do not select the Thin Provisioning option, your virtual disk will have the default thick format.

What to do next

If you created a virtual disk in the thin format, you can later inflate it to its full size.

View Virtual Machine Storage Resources

You can view how datastore storage space is allocated for your virtual machines.

Procedure

- 1 Select the virtual machine in the inventory.
- 2 Click the **Summary** tab.

- 3 Review the space allocation information in the Resources section.
 - **Provisioned Storage** – Shows datastore space guaranteed to the virtual machine. The entire space might not be used by the virtual machine if it has disks in thin provisioned format. Other virtual machines can occupy any unused space.
 - **Not-shared Storage** – Shows datastore space occupied by the virtual machine and not shared with any other virtual machines.
 - **Used Storage** – Shows datastore space actually occupied by virtual machine files, including configuration and log files, snapshots, virtual disks, and so on. When the virtual machine is running, the used storage space also includes swap files.

Determine the Disk Format of a Virtual Machine

You can determine whether your virtual disk is in thick or thin format.

Procedure

- 1 Select the virtual machine in the inventory.
- 2 Click **Edit Settings** to display the Virtual Machine Properties dialog box.
- 3 Click the **Hardware** tab and select the appropriate hard disk in the Hardware list.

The Disk Provisioning section on the right shows the type of your virtual disk, either Thin or Thick.

- 4 Click **OK**.

What to do next

If your virtual disk is in the thin format, you can inflate it to its full size.

Convert a Virtual Disk from Thin to Thick

If you created a virtual disk in the thin format, you can convert it to thick.

Procedure

- 1 Select the virtual machine in the inventory.
- 2 Click the **Summary** tab and, under Resources, double-click the datastore for the virtual machine to open the Datastore Browser dialog box.
- 3 Click the virtual machine folder to find the virtual disk file you want to convert. The file has the `.vmdk` extension.
- 4 Right-click the virtual disk file and select **Inflate**.

The virtual disk in thick format occupies the entire datastore space originally provisioned to it.

Handling Datastore Over-Subscription

Because the provisioned space for thin disks can be greater than the committed space, a datastore over-subscription can occur, which results in the total provisioned space for the virtual machine disks on the datastore being greater than the actual capacity.

Over-subscription can be possible because usually not all virtual machines with thin disks need the entire provisioned datastore space simultaneously. However, if you want to avoid over-subscribing the datastore, you can set up an alarm that notifies you when the provisioned space reaches a certain threshold.

For information on setting alarms, see *Basic System Administration*.

If your virtual machines require more space, the datastore space is allocated on a first come first served basis. When the datastore runs out of space, you can add more physical storage and increase the datastore.

See [“Increase VMFS Datastores,”](#) on page 101.

TurnOff vCenter Server Storage Filters

When you perform VMFS datastore management operations, for example create a VMFS datastore or an RDM, add an extent, or increase a VMFS datastore, the vCenter Server uses default storage filters. The filters retrieve only those storage devices, or LUNs, that can be used for a particular operation. Unsuitable LUNs are filtered out and are not displayed for your selection. Using the vSphere Client, you can turn off the filters.

To turn off the filters, you use the Advanced Settings page to modify the vCenter Server configuration file, `vpxd.cfg`. This page can be used to add entries to the file, but not to edit or delete them.

As with any advanced configuration settings, consult with the VMware support team before making any changes to the LUN filters. Do not turn the filters off unless you have other methods to prevent LUN corruption.

Procedure

- 1 Select **Administration > vCenter Server Settings** to display the vCenter Server Settings dialog box.
- 2 In the settings list, select **Advanced Settings**.
- 3 In the **Key** field, type a key.

Key	Description
<code>config.vpxd.filter.vmfsFilter</code>	VMFS Filter
<code>config.vpxd.filter.rdmFilter</code>	RDM Filter
<code>config.vpxd.filter.SameHostAndTransportsFilter</code>	Same Host and Transports Filter
<code>config.vpxd.filter.hostRescanFilter</code>	Host Rescan Filter

- 4 In the **Value** field, type `False` for the specified key.
- 5 Click **Add**.
- 6 Click **OK**.

What to do next

You are not required to restart the vCenter Server system.

vCenter Server Storage Filtering

The storage filters that the vCenter Server provides help you avoid storage device corruption and performance degradation that can be caused by an unsupported use of LUNs. These filters are available by default.

Table 9-2. Storage Filters

Filter	Description	Key
VMFS Filter	Filters out any storage devices, or LUNs, that are already used by another VMFS datastore on any host managed by the vCenter Server. Prevents LUN sharing by multiple datastores or a datastore and RDM combination.	config.vpxd.filter.vmfsFilter
RDM Filter	Filters out any LUNs that are already referenced by another RDM on any host managed by the vCenter Server. Prevents LUN sharing by a datastore and RDM combination. In addition, the filter prevents virtual machines from accessing the same LUN through different RDM mapping files. If you need virtual machines to access the same raw LUN, they must share the same RDM mapping file. For details on this type of configuration, see <i>Setup for Failover Clustering and Microsoft Cluster Service</i> .	config.vpxd.filter.rdmFilter
Same Host and Transports Filter	Filters out LUNs ineligible for use as VMFS datastore extents due a host or storage type incompatibility. Prevents you from adding the following LUNs as extents: ■ LUNs not exposed to all hosts that share the original VMFS datastore. ■ LUNs that use a storage type different from the one the original VMFS datastore uses. For example, you cannot add a Fibre Channel extent to a VMFS datastore on a local storage device.	config.vpxd.filter.SameHostAndTransportsFilter
Host Rescan Filter	Automatically rescans and updates storage devices after you perform datastore management operations. The filter helps provide a consistent view of all storage devices and VMFS datastores on all hosts managed by the vCenter Server.	config.vpxd.filter.hostRescanFilter

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).

The following topics contain information about RDMs and provide instructions on how to create and manage RDMs.

This chapter includes the following topics:

- [“About Raw Device Mapping,”](#) on page 119
- [“Raw Device Mapping Characteristics,”](#) on page 122
- [“Managing Mapped LUNs,”](#) on page 126


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 for managing and redirecting disk access 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 VMFS. 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 10-1. 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. 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 vSphere 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.

RDM offers several benefits.

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. VMFS 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. 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 vCenter Server to migrate the virtual machine by using the same mechanism that exists for migrating virtual disk files.

Figure 10-2. 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. 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 You can use NPIV only for virtual machines with RDM disks.

VMware works with vendors of storage management software to ensure that their software functions correctly in environments that include ESX. 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 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 environment, contact the SAN management software provider.

Limitations of Raw Device Mapping

Certain limitations exist when you use RDMs.

- Not available for block devices or certain RAID devices – RDM 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, 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 in virtual mode.

- 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), the locking state of the mapped device, permissions, and so on.

RDM Virtual and Physical Compatibility Modes

You can use RDMs in virtual compatibility or physical compatibility modes. Virtual mode specifies full virtualization of the mapped device. Physical mode specifies minimal SCSI virtualization of the mapped device, allowing the greatest flexibility for SAN management software.

In virtual mode, the mapped device appears to the guest operating system exactly the same as a virtual disk file in a VMFS volume. The real hardware characteristics are hidden. If you are using a raw disk in virtual mode, you can 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.


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 10-3. 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 10-4](#) 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 10-4. 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 device name. 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 10-5. Access from Clustered Virtual Machines



Comparing Available SCSI Device Access Modes

The ways of accessing a SCSI-based storage device include a virtual disk file on a VMFS datastore, virtual mode RDM, and physical mode RDM.

To help you choose among the available access modes for SCSI devices, [Table 10-1](#) provides a quick comparison of features available with the different modes.

Table 10-1. Features Available with Virtual Disks and Raw Device Mappings

ESX Features	Virtual Disk File	Virtual Mode RDM	Physical Mode RDM
SCSI Commands Passed Through	No	No	Yes REPORT LUNs is not passed through
vCenter Server Support	Yes	Yes	Yes
Snapshots	Yes	Yes	No
Distributed Locking	Yes	Yes	Yes
Clustering	Cluster-in-a-box only	Cluster-in-a-box and cluster-across-boxes	Physical to Virtual Clustering
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.

Managing Mapped LUNs

You can use the vSphere Client to map a SAN LUN to a datastore and manage paths to your mapped LUN.

Additional tools available to manage mapped LUNs and their RDMs include the `vmkfstools` utility and other commands used with the vSphere CLI. You can use the `vmkfstools` utility to perform many of the same operations available through the vSphere Client.

You can also use common file system commands in the service console.

Create 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.

Procedure

- 1 Follow all steps required to create a custom virtual machine.
- 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.
- 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.

Option	Description
Physical	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 a physical compatibility RDM cannot be cloned, made into a template, or migrated if the migration involves copying the disk.
Virtual	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.

Option	Description
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.

Manage Paths for a Mapped Raw LUN

You can manage paths for mapped raw LUNs.

Procedure

- 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 **Edit Settings**.
- 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.

For information on managing paths, see [“Using Multipathing with ESX,”](#) on page 105.

Security

Security for ESX Systems

ESX is developed with a focus on strong security. VMware ensures security in the ESX environment and addresses system architecture from a security standpoint.

This chapter includes the following topics:

- [“ESX Architecture and Security Features,”](#) on page 131
- [“Security Resources and Information,”](#) on page 139


ESX Architecture and Security Features

The components and the overall architecture of ESX are designed to ensure security of the ESX system as a whole.

From a security perspective, ESX consists of four major components: the virtualization layer, the virtual machines, the service console, and the virtual networking layer.

[Figure 11-1](#) provides an overview of these components.

Figure 11-1. ESX Architecture



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 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.

ESX provides additional VMkernel protection with the following features:

Memory Hardening	The ESX kernel, user-mode applications, and executable components such as drivers and libraries are located at random, non-predictable memory addresses. Combined with the non-executable memory protections made available by microprocessors, this provides protection that makes it difficult for malicious code to use memory exploits to take advantage of vulnerabilities.
Kernel Module Integrity	Digital signing ensures the integrity and authenticity of modules, drivers and applications as they are loaded by the VMkernel. Module signing allows ESX to identify the providers of modules, drivers, or applications and whether they are VMware-certified.


Security and Virtual Machines

Virtual machines are the containers in which applications and guest operating systems run. By design, all VMware virtual machines are isolated from one another. This isolation enables multiple virtual machines to run securely while sharing hardware and ensures both their ability to access hardware and their uninterrupted performance.

Even a user with system administrator privileges on a virtual machine's guest operating system cannot breach this layer of isolation to access another virtual machine without privileges explicitly granted by the ESX system administrator. As a result of virtual machine isolation, if a guest operating system running in a virtual machine fails, other virtual machines on the same host continue to run. The guest operating system failure 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. Although 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 11-2](#).

Figure 11-2. 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 communicates with other machines in a network through a network card, a virtual machine communicates with other virtual machines running in the same host through a virtual switch. Further, a virtual machine communicates with the physical network, including virtual machines on other ESX hosts, through a physical network adapter, as shown in [Figure 11-3](#).

Figure 11-3. Virtual Networking Through Virtual Switches

These characteristics apply to virtual machine isolation in a network context:

- 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.

You can further protect virtual machines by setting up resource reservations and limits on the host. For example, through the detailed resource controls available in ESX, you can configure a virtual machine so that it always receives at least 10 percent of the host's CPU resources, but never more than 20 percent.

Resource reservations and limits protect virtual machines from performance degradation that would result if another virtual machine consumed excessive shared hardware resources. For example, if one of the virtual machines on a host is incapacitated by a denial-of-service (DoS) attack, a resource limit on that machine prevents the attack from taking up so much of the 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 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. This default behavior provides a degree of natural protection from DoS and distributed denial-of-service (DDoS) attacks. You set specific resource reservations and limits on an individual basis to customize the default behavior so that the distribution is not equal across the virtual machine configuration.

Security and the Virtual Networking Layer

The virtual networking layer includes virtual network adapters and virtual switches. ESX relies on the virtual networking layer to support communications between virtual machines and their users. In addition, hosts use the virtual networking layer to communicate with iSCSI SANs, NAS storage, and so forth.


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 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.

Creating a Network DMZ on a Single ESX Host

One example of how to use ESX isolation and virtual networking features to configure a secure environment is the creation of a network demilitarized zone (DMZ) on a single host.

[Figure 11-4](#) shows the configuration.

Figure 11-4. DMZ Configured on a Single ESX Host

In this example, four virtual machines are 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 multi-homed.
- 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 host, 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. This network could be used for virus propagation or targeted for other types of attacks. The security of the virtual machines in the DMZ is equivalent to separate physical machines connected to the same network.

Creating Multiple Networks Within a Single ESX Host

The ESX 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 host. This capability is an outgrowth of basic virtual machine isolation coupled with a well-planned use of virtual networking features.

Figure 11-5. External Networks, Internal Networks, and a DMZ Configured on a Single ESX Host



In [Figure 11-5](#) the system administrator configured a host into three distinct virtual machine zones: FTP server, internal virtual machines, and DMZ. Each zone serves 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 hosts throughout the site.

Because Virtual Machine 1 does not 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 adjusters.

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 content 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.

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 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 under the following circumstances:

- The virtual switches are connected to the same physical LAN.
- The virtual switches connect to a common virtual machine, which could 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 vSphere Client or vSphere Web Access.

To safeguard the virtual machines' resources, the system administrator lowers the risk of DoS and DDoS attacks by configuring a resource reservation and a limit for each virtual machine. The system administrator further protects the ESX host and virtual machines by installing software firewalls at the front and back ends of the DMZ, ensuring that the host is behind a physical firewall, and configuring the service console and networked storage resources so that each has its own virtual switch.

Security and the Service Console

The ESX service console is a limited distribution of Linux based on Red Hat Enterprise Linux 5 (RHEL5). The service console provides an execution environment to monitor and administer the entire ESX 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.

In addition to implementing the service console firewall, VMware mitigates risks to the service console using other methods.

- ESX runs only services essential to managing its functions, and the distribution is limited to the features required to run ESX.
- By default, ESX is installed with a high-security setting. All outbound ports are closed, and the only inbound ports that are open are those required for interactions with clients such as the vSphere Client. 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.
- By default, weak ciphers are disabled and all communications from clients are secured by SSL. The exact algorithms used for securing the channel depend on the SSL handshake. Default certificates created on ESX use SHA-1 with RSA encryption as the signature algorithm.
- The Tomcat Web service, used internally by ESX to support access to the service console by Web clients like vSphere Web Access, has been modified to run only those functions required for administration and monitoring by a Web client. As a result, ESX is not vulnerable to the Tomcat security issues reported in broader use.
- 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 hosts. VMware provides security patches for RHEL 5 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. You can disable any `setuid` or `setgid` application that is optional to ESX operation.

Although you can install and run certain types of programs designed for RHEL 5 in the service console, this use 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 Follow only VMware security advisories, found at <http://www.vmware.com/security/>. Do not follow security advisories issued by Red Hat.

Security Resources and Information

You can find additional information about security on the VMware Web site.

[Table 11-1](#) lists security topics and the location of additional information about these topics.

Table 11-1. VMware Security Resources on the Web

Topic	Resource
VMware security policy, up-to-date security alerts, security downloads, and focus discussions of security topics	http://www.vmware.com/security/
Corporate security response policy	http://www.vmware.com/support/policies/security_response.html VMware is committed to helping you maintain a secure environment. Security issues are corrected in a timely manner. 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 by searching http://www.vmware.com/vmtn/resources/ for ESX 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 security risks for unsupported products or configurations carefully.
Certification of VMware products	http://www.vmware.com/security/certifications/

Securing an ESX Configuration

You can take measures to promote a secure environment for your ESX hosts, virtual machines, and iSCSI SANs. Consider network configuration planning from a security perspective and the steps that you can take to protect the components in your configuration from attack.

This chapter includes the following topics:

- [“Securing the Network with Firewalls,”](#) on page 141
- [“Securing Virtual Machines with VLANs,”](#) on page 150
- [“Securing Virtual Switch Ports,”](#) on page 155
- [“Securing iSCSI Storage,”](#) on page 157

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. 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 components.

- Physical machines such as vCenter Server hosts and ESX 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 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. Therefore, a configuration where firewalls are present between the virtual machines is not necessary. 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.

Firewalls for Configurations with vCenter Server


If you access ESX hosts through vCenter Server, you typically protect vCenter Server using a firewall. This firewall provides basic protection for your network.

A firewall might lie between the clients and vCenter Server. Alternatively, vCenter Server and the clients can be behind the firewall, depending 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.

If you use vCenter Server, you can install firewalls at any of the locations shown in [Figure 12-1](#). Depending on your configuration, you might not need all the firewalls in the illustration, or you might need firewalls in other locations. In addition, your configuration might include optional modules, such as VMware vCenter Update Manager, that are not shown. Refer to the documentation for information about firewall setups specific to products like Update Manager.

For a comprehensive list of TCP and UDP ports, including those for VMware VMotion™ and VMware Fault Tolerance, see [“TCP and UDP Ports for Management Access,”](#) on page 149.

Figure 12-1. Sample vSphere Network Configuration and Traffic Flow



Networks configured with vCenter Server can receive communications through several types of clients: the vSphere Client, vSphere Web Access, or third-party network management clients that use the SDK to interface with the host. During normal operation, vCenter Server listens for data from its managed hosts and clients on designated ports. vCenter Server also assumes that its managed hosts listen for data from vCenter Server 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.

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 the level of security various devices require. Select the locations for your firewalls based on the security risks that you have identified for your network configuration. The following is a list of firewall locations common to ESX implementations. Many of the firewall locations in the list and shown in [Figure 12-1](#) are optional.

- Between your Web browser and the vSphere Web Access HTTP and HTTPS proxy server.
- Between the vSphere Client, vSphere Web Access Client, or a third-party network-management client and vCenter Server.
- If your users access virtual machines through the vSphere Client, between the vSphere Client and the ESX host. This connection is in addition to the connection between the vSphere Client and vCenter Server, and it requires a different port.
- If your users access virtual machines through a Web browser, between the Web browser and the ESX host. This connection is in addition to the connection between the vSphere Web Access Client and vCenter Server, and it requires different ports.
- Between vCenter Server and the ESX hosts.
- Between the ESX hosts in your network. Although traffic between hosts is usually considered trusted, you can add firewalls between them if you are concerned about security breaches from machine to machine.

If you add firewalls between ESX 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 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.

Firewalls for Configurations Without vCenter Server

If you connect clients directly to your ESX network instead of using vCenter Server, your firewall configuration is somewhat simpler.

You might install firewalls at any of the locations shown in [Figure 12-2](#).

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 12-2. Firewall Configuration for ESX Networks that a Client Manages Directly

Networks configured without vCenter Server receive communications through the same types of clients as they do if vCenter Server were present: vSphere Clients, third-party network management clients, or vSphere Web Access Clients. For the most part, the firewall needs are the same, but there are several key differences.

- As you would for configurations that include vCenter Server, be sure a firewall is present to protect your ESX layer or, depending on your configuration, your clients and ESX layer. This firewall provides basic protection for your network. The firewall ports you use are the same as those you use if vCenter Server is in place.
- Licensing in this type of configuration is part of the ESX package that you install on each of the hosts. Because licensing is resident to the server, a separate license server is not required. This eliminates the need for a firewall between the license server and the ESX network.

Connecting to vCenter Server Through a Firewall

The port that vCenter Server uses to listen for data transfer from its clients is 443. If you have a firewall between vCenter Server and its clients, you must configure a connection through which vCenter Server can receive data from the clients.

To enable vCenter Server to receive data from the vSphere Client, open port 443 in the firewall to allow data transfer from the vSphere Client to vCenter Server. Contact the firewall system administrator for additional information on configuring ports in a firewall.

If you are using the vSphere Client and do not want to use port 443 as the port for vSphere Client-to-vCenter Server communication, you can switch to another port by changing the vCenter Server settings in the vSphere 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 hosts through vCenter Server or use a direct connection to the 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, and use different authentication protocols.

Port 902

vCenter Server uses this port to send data to vCenter Server managed hosts. Port 902 is the port that vCenter Server assumes is available when sending data to an ESX host.

Port 902 connects vCenter Server to the host through the VMware Authorization Daemon (`vmware-authd`). This daemon multiplexes port 902 data to the appropriate recipient for processing. VMware does not support configuring a different port for this connection.

Port 443

The vSphere Client, vSphere Web Access Client, and SDK use this port to send data to vCenter Server managed hosts. Also, the vSphere Client, vSphere Web Access Client, and SDK, when connected directly to an ESX 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 host. VMware does not support configuring a different port for these connections.

Port 443 connects clients to the ESX host through the Tomcat Web service or the SDK. The `vmware-hostd` multiplexes port 443 data to the appropriate recipient for processing.

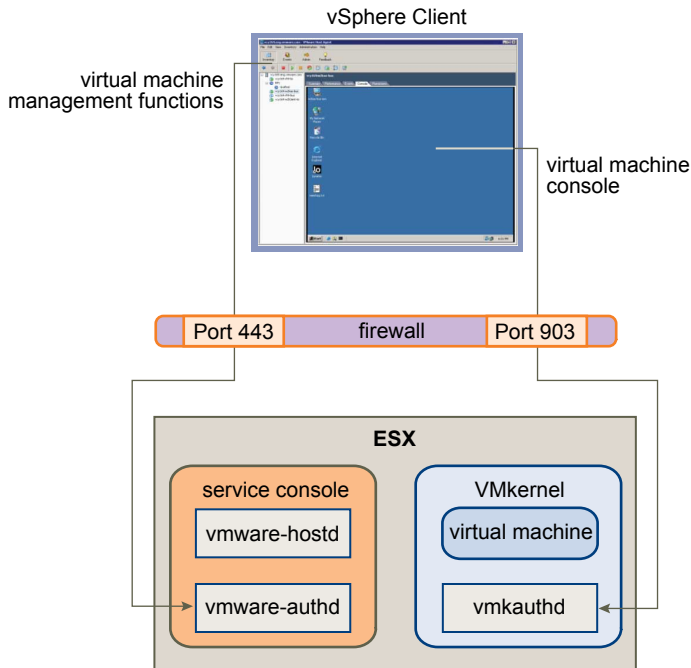
Port 903

The vSphere Client and vSphere Web Access use this port to provide a connection for guest operating system 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 vSphere Client and vSphere 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 vSphere Client to a specified virtual machine configured on the ESX host.

[Figure 12-3](#) shows the relationships between vSphere Client functions, ports, and ESX processes.

The vSphere Web Access Client uses the same basic mapping for its interactions with the ESX host.

Figure 12-3. Port Use for vSphere Client Communications with ESX

If you have a firewall between your vCenter Server system and vCenter Server managed host, open Ports 443 and 903 in the firewall to allow data transfer to ESX hosts from vCenter Server and ESX hosts directly from the vSphere Client and vSphere Web Access.

For additional information on configuring the ports, see the firewall system administrator.

Connecting ESX Hosts Through Firewalls

If you have a firewall between two ESX hosts and you want to allow transactions between the hosts or use vCenter Server 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 configure a connection for receiving data, 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)

Refer to the firewall system administrator for additional information on configuring the ports.

Configuring Firewall Ports for Supported Services and Management Agents

You must configure firewalls in your environment to accept commonly supported services and installed management agents.

Use the vSphere Client to configure the service console firewall. When you configure the ESX host security profile in vCenter Server, 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 services and agents are commonly present in a vSphere environment:

- NFS client (insecure service)

- NTP client
- iSCSI software client
- CIM HTTP server (insecure service)
- CIM HTTPS server
- Syslog client
- NFS server (insecure service)
- NIS client
- SMB client (insecure service)
- FTP client (insecure service)
- SSH client
- Telnet client (insecure service)
- SSH server
- Telnet server (insecure service)
- FTP server (insecure service)
- SNMP server
- Other supported management agents that you install

NOTE This list can change, so you might find that the vSphere Client provides services and agents not mentioned in the list. Also, not all services on the list are installed by default. You might be required 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.

Allow Access to ESX for a Service or Management Agent

You can configure firewall properties to allow access for a service or management agent.

Procedure

- 1 Log in to a vCenter Server system using the vSphere Client.
- 2 Select the host in the inventory panel.
- 3 Click the **Configuration** tab and click **Security Profile**.

The vSphere Client displays a list of active incoming and outgoing connections with the corresponding firewall ports.

- 4 Click **Properties** to open the Firewall Properties dialog box.

The Firewall Properties dialog box lists all the services and management agents that you can configure for the host.

- 5 Select the services and agents to enable.

The Incoming Ports and Outgoing Ports columns indicate the ports that the vSphere Client opens for the service. The Protocol column indicates the protocol that the service uses. The Daemon column indicates the status of daemons associated with the service.

- 6 Click **OK**.

Automating Service Behavior Based on Firewall Settings

ESX can automate whether services start based on the status of firewall ports.

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.

In addition, 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. 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 vSphere Client or applications created with the vSphere Web services 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. 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. 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 shutdown process, but as soon as the host is powered on, the service is started again, preserving the user-determined state.

Configure How Service Startup Relates to Firewall Configuration

The Startup Policy determines when a service starts. You can configure how service startup relates to a firewall configuration by editing the Startup Policy.

Procedure

- 1 Log in to a vCenter Server system using the vSphere Client.
- 2 Select the host in the inventory panel.
- 3 Click the **Configuration** tab and click **Security Profile**.

The vSphere Client displays a list of active incoming and outgoing connections with the corresponding firewall ports.

- 4 Click **Properties**.

The Firewall Properties dialog box lists all the services and management agents you can configure for the host.

- 5 Select the service to configure, and click **Options**.

The Startup Policy dialog box 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.

- 6 Select a policy from the **Startup Policy** list.
- 7 Click **OK**.

TCP and UDP Ports for Management Access

vCenter Server, ESX hosts, and other network components are accessed using predetermined TCP and UDP ports. If you manage network components from outside a firewall, you might be required to reconfigure the firewall to allow access on the appropriate ports.

[Table 12-1](#) lists TCP and UDP ports, and the purpose and the type of each.

The ports are connected through the service console interface, unless otherwise indicated.

Table 12-1. TCP and UDP Ports

Port	Purpose	Traffic Type
22	SSH Server	Incoming TCP
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 networks from the Web. Port 80 redirects traffic to an HTTPS landing page (port 443). Connection to vSphere Web Access from the Web WS-Management	Incoming TCP
123	NTP Client	Outgoing UDP
427	The CIM client uses the Service Location Protocol, version 2 (SLPv2) to find CIM servers.	Incoming and outgoing UDP
443	HTTPS access vCenter Server access to ESX hosts Default SSL Web port vSphere Client access to vCenter Server vSphere Client access to ESX hosts WS-Management vSphere Client access to vSphere Update Manager vSphere Converter access to vCenter Server vSphere Web Access and third-party network management client connections to vCenter Server Direct vSphere Web Access and third-party network management clients access to hosts	Incoming TCP
902	Host access to other hosts for migration and provisioning Authentication traffic for ESX (xinetd/vmware-authd) vSphere Client access to virtual machine consoles (UDP) Status update (heartbeat) connection from ESX to vCenter Server	Incoming TCP, outgoing UDP
903	Remote console traffic that user access to virtual machines generates on a specific ESX host. vSphere Client access to virtual machine consoles vSphere Web Access Client access to virtual machine consoles MKS transactions (xinetd/vmware-authd-mks)	Incoming TCP
2049	Transactions from NFS storage devices This port is used on the VMkernel interface rather than the service console interface.	Incoming and outgoing TCP

Table 12-1. TCP and UDP Ports (Continued)

Port	Purpose	Traffic Type
2050–2250	Traffic between ESX hosts for VMware High Availability (HA) and EMC Autostart Manager	Outgoing TCP, incoming and outgoing UDP
3260	Transactions to iSCSI storage devices This port is used on the VMkernel interface and the service console interface.	Outgoing TCP
5900-5964	RFB protocol, which is used by management tools such as VNC	Incoming and outgoing TCP
5989	CIM XML transactions over HTTPS	Incoming and outgoing TCP
8000	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 hosts for HA and EMC Autostart Manager	Outgoing TCP, incoming and outgoing UDP
8100, 8200	Traffic between ESX hosts for VMware Fault Tolerance	Outgoing TCP, incoming and outgoing UDP

In addition to the TCP and UDP ports listed in [Table 12-1](#), you can configure other ports depending on your needs:

- You can use vSphere Client to open ports for installed management agents and supported services such as NFS.
- You can open ports in the service console firewall for other services and agents required for your network by running command-line scripts.

Securing Virtual Machines with VLANs

The network can be one of the most vulnerable parts of any system. Your virtual machine network requires as much protection as your physical network. You can add security to your virtual machine network in several ways.

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 is. 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 in different ways.

- 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.


- 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, which prevents sniffing attacks that 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 12-4](#).

Figure 12-4. 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 VLANs

The way you set up VLANs to secure parts of a network depends on factors such as the guest operating system and the way your network equipment is configured.

ESX features a complete IEEE 802.1q-compliant VLAN implementation. VMware cannot make specific recommendations on how to set up VLANs, but there are factors to consider when using a VLAN deployment as part of your security enforcement policy.

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 does not 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 the firewalls that you install are VLAN-aware.

Properly Configure VLANs

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, disable the native VLAN feature for all switches unless you have a compelling reason to operate some of your VLANs in native mode. If you must use native VLAN, see your switch vendor's configuration guidelines for this feature.

Create Separate Communications Between Management Tools and the Service Console

Whether you use a management client or the command line, all configuration tasks for ESX 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, safeguarding it from misuse is crucial.

VMware ESX 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 host. For example, they can change the entire virtual switch configuration or change authorization methods.

Network connectivity for the service console is established through virtual switches. To provide better protection for this critical ESX component, 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.

Set up a separate VLAN or virtual switch for VMotion and network attached storage.

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; they protect only the virtual network.

Virtual switches and VLANs can protect against the following types of attacks.

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.

Although VMware virtual switches store a MAC address table, they do not 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 do not 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 to overload 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 do not 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 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 and storage systems attached to the switch. You create virtual ports through the vSphere Client.

As part of adding a port or port group to a virtual switch, the vSphere Client configures a security profile for the port. You can use this security profile to ensure that ESX 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 must 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 receives network traffic destined for the new MAC address. The operating system can send frames with an impersonated source MAC address at any time. This means 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 hosts to protect against this type of attack by setting three options. If you change any default settings for a port, you must modify the security profile by editing virtual switch settings in the vSphere Client.

MAC Address Changes

The setting for the **MAC Address Changes** option affects traffic that a virtual machine receives.

When the option is set to **Accept**, ESX accepts requests to change the effective MAC address to other than the initial MAC address.

When the option is set to **Reject**, ESX does not honor requests to change the effective MAC address to anything other than the initial MAC address, which protects the host against MAC impersonation. The port that the virtual adapter used to send the request is disabled and 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 The iSCSI initiator relies on being able to get MAC address changes from certain types of storage. If you are using ESX iSCSI and have iSCSI storage, set the **MAC Address Changes** option to **Accept**.

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

The setting for the **Forged Transmits** option affects traffic that is transmitted from a virtual machine.

When the option is set to **Accept**, ESX does not compare source and effective MAC addresses.

To protect against MAC impersonation, you can set this option to **Reject**. If you do, the 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 do not match, ESX 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 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

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. By default, the virtual network adapter cannot operate in promiscuous mode.

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 are 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 reason to configure a virtual switch to operate in promiscuous mode—for example, if you are running network intrusion detection software or a packet sniffer.

Securing iSCSI Storage

The storage you configure for an ESX host might include one or more storage area networks (SANs) that use iSCSI. When you configure iSCSI on an ESX host, you can take several measures to minimize security risks.

iSCSI is a means of accessing SCSI devices and exchanging data records by using TCP/IP 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 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.

NOTE The requirements and procedures for securing an iSCSI SAN are similar for the hardware iSCSI adapters you can use with ESX hosts and for iSCSI configured directly through the ESX host.

Securing iSCSI Devices Through Authentication

One means of securing iSCSI devices from unwanted intrusion is to require that the ESX 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.

ESX 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 vSphere Client to determine whether authentication is being performed and to configure the authentication method.

Enabling Challenge Handshake Authentication Protocol (CHAP) for iSCSI SANs

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 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 supports unidirectional and bidirectional CHAP authentication for iSCSI. In unidirectional CHAP authentication, the target authenticates the initiator, but the initiator does not authenticate the target. In bidirectional CHAP authentication, an additional level of security enables the initiator to authenticate the target.

ESX supports CHAP authentication at the adapter level, when only one set of authentication credentials can be sent from the host to all targets. It also supports per-target CHAP authentication, which enables you to configure different credentials for each target to achieve greater target refinement.

See [“Configuring CHAP Parameters for iSCSI Initiators,”](#) on page 89 for information about how to work with CHAP.

Disabling iSCSI SAN Authentication

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 is.

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 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 vSphere Client to determine whether authentication is being performed and to configure the authentication method.

See [“Configuring CHAP Parameters for iSCSI Initiators,”](#) on page 89 for information about how to work with CHAP.

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 for enforcing good security standards.


Protect Transmitted Data

A primary security risk in iSCSI SANs is that an attacker might sniff transmitted storage data.

Take additional measures to prevent attackers from easily seeing iSCSI data. Neither the hardware iSCSI adapter nor the ESX host iSCSI initiator encrypts the data that 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 cannot 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 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 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 12-5](#).

Figure 12-5. iSCSI Storage on a Separate Virtual Switch

In addition to protecting the iSCSI SAN by giving it a dedicated virtual switch, you can configure your iSCSI SAN on its own VLAN to improve performance and security. 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. Also, network congestion from other sources cannot interfere with iSCSI traffic.

Secure iSCSI Ports

When you run iSCSI devices, the ESX host does not open any ports that listen for network connections. This measure reduces the chances that an intruder can break into the ESX host through spare ports and gain control over the host. Therefore, running iSCSI does not present any additional security risks at the ESX host end of the connection.

Any iSCSI target device that you run must have one or more open TCP ports 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. 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

ESX handles user authentication and supports user and group permissions. In addition, you can encrypt connections to the vSphere Client and SDK.

This chapter includes the following topics:

- [“Securing ESX Through Authentication and Permissions,”](#) on page 161
- [“Encryption and Security Certificates for ESX,”](#) on page 168

Securing ESX Through Authentication and Permissions

When a vSphere Client or vCenter Server user connects to a ESX host, a connection is established with the VMware Host Agent process. The process uses the user names and passwords for authentication.


ESX uses the Pluggable Authentication Modules (PAM) structure for authentication when users access the ESX host using the vSphere Client, vSphere 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 uses `/etc/passwd` authentication as Linux does, but you can configure ESX to use another distributed authentication mechanism. If you plan to use a third-party authentication tool instead of the ESX default implementation, see the vendor documentation for instructions. As part of setting up third-party authentication, you might be required to update the files in `/etc/pam.d` folder with new module information.

The reverse proxy in the VMware Host Agent (`vmware-hostd`) process listens on ports 80 and 443. vSphere Client or vCenter Server users connect to the host agent through these ports. 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 13-1](#) shows a basic example of how ESX authenticates transactions from the vSphere Client.

NOTE CIM transactions also use ticket-based authentication in connecting with the `vmware-hostd` process.

Figure 13-1. Authentication for vSphere Client Communications with ESX

ESX authentication transactions with vSphere Web Access and third-party network management clients are also direct interactions with the `vmware-hostd` process.

To make sure that authentication works efficiently for your site, perform basic tasks such as setting up users, groups, permissions, and roles, configuring user attributes, adding your own certificates, and determining whether you want to use SSL.

About Users, Groups, Permissions, and Roles

vCenter Server and ESX hosts use a combination of user name, password, and permissions to authenticate a user for access and authorize activities. You can control access to hosts, clusters, datastores, resource pools, networking port groups, and virtual machines by assigning permissions.

Access to an ESX host and its resources is granted when a known user with appropriate permissions logs in to the host with a correct password. vCenter Server uses a similar approach when determining whether to grant access to a user.

vCenter Server and ESX 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 hosts and vCenter Server, you must plan how to handle particular types of users and permissions. ESX and vCenter Server use sets of privileges, or roles, to control which operations individual users or groups can perform. Predefined roles are provided, but you can also create new ones. You can manage users more easily by assigning them to groups. When you apply a role to the group, all users in the group inherit the role.

Understanding Users

A user is an individual authorized to log in to either an ESX host or vCenter Server.

ESX users fall into two categories: those who can access the host through vCenter Server and those who can access by directly logging in to the host from the vSphere Client, vSphere Web Access, a third-party client, or a command shell.

Authorized vCenter Server users

Authorized users for vCenter Server are those included in the Windows domain list that vCenter Server references or are local Windows users on the vCenter Server host.

You cannot use vCenter Server to manually create, remove, or otherwise change users. You must use the tools for managing your Windows domain. Any changes you make are reflected in vCenter Server. However, the user interface does not provide a user list for you to review.

Direct-access users

Users authorized to work directly on an ESX host are those added to the internal user list by a system administrator.

An administrator can perform a variety of management activities for these users, such as changing passwords, group memberships, and permissions as well as adding and removing users.

The user list that vCenter Server maintains is separate from the user list that the host maintains. Even if the lists appear to have common users (for instance, a user called devuser), treat these users separately. If you log in to vCenter Server as devuser, you might have permission to view and delete files from a datastore, whereas if you log in to an ESX host as devuser, you might not.

Because of the confusion that duplicate naming can cause, check the vCenter Server user list before you create ESX host users to avoid duplicating names. To check for vCenter Server users, review the Windows domain list.

Understanding Groups

A group is a set of users that share 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.

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 and give the group permission to interact with only one object, the shared virtual machine. They cannot perform any actions on the sales managers' virtual machines.

The group lists in vCenter Server and an ESX host are drawn from the same sources as their respective user lists. If you are working through vCenter Server, the group list is called from the Windows domain. If you are logged in to an ESX host directly, the group list is called from a table that the host maintains.

Understanding Permissions

For ESX and vCenter Server, 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 host.

Most vCenter Server and ESX users have limited ability to manipulate the objects associated with the host. Users with 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 vCenter Server manages the host, vpxuser is also an Administrator user.

The list of privileges is the same for both ESX and vCenter Server, and you use the same method to configure permissions.

You can create roles and set permissions through a direct connection to the ESX host. Because these tasks are widely performed in vCenter Server, see *Basic System Administration* for information on working with permissions and roles.

Assigning root User Permissions

Root users can only perform activities on the specific ESX host that they are logged in to.

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 vSphere 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 vSphere Client logs all actions that the Administrator role user initiates as events, providing you with an audit trail. If all administrators log in as the root user, you cannot tell which administrator performed an action. If 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 delete the root user's permissions or change the role to limit its privileges. You must then use the new user you created as the host authentication point when you bring the host under vCenter Server management.

NOTE `vicfg` commands do not perform an access check. Therefore, even if you limit the root user's privileges, it does not affect what that user can do using the command-line interface commands.

Understanding vpxuser Permissions

The vpxuser permission is used for vCenter Server when managing activities for the host. The vpxuser is created when an ESX host is attached to vCenter Server.

vCenter Server has Administrator privileges on the host that it manages. For example, vCenter Server can move virtual machines to and from hosts and perform configuration changes needed to support virtual machines.

The vCenter Server administrator 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, the vCenter Server administrator cannot directly create, delete, or edit users and groups for ESX hosts. These tasks can only be performed by a user with Administrator permissions directly on each ESX 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 ESX hosts through vCenter Server.

Understanding Roles

vCenter Server and ESX 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 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.

You can create custom roles by using the role-editing facilities in the vSphere Client to create privilege sets that match your user needs. If you use the vSphere Client connected to vCenter Server to manage your ESX hosts, you have additional roles to choose from in vCenter Server. Also, the roles you create directly on an ESX host are not accessible within vCenter Server. You can work with these roles only if you log in to the host directly from the vSphere Client.

If you manage ESX hosts through vCenter Server, maintaining custom roles in the host and vCenter Server can result in confusion and misuse. In this type of configuration, maintain custom roles only in vCenter Server.

You can create roles and set permissions through a direct connection to the ESX host. Because most users create roles and set permissions in vCenter Server, see *Basic System Administration* for information on working with permissions and roles.

Assigning the No Access Role

Users assigned the No Access role for an object cannot view or change the object in any way. New users and groups are assigned this role by default. You can change the role on an object-by-object basis.

A user with a No Access role for a particular object can select the vSphere Client tabs associated with the object, but the tab displays no content.

The root user and vpxuser permissions 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.

Assigning the Read Only Role

Users assigned the Read Only 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.

Assigning the Administrator Role

Users assigned the Administrator 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.

If you are acting in the Administrator role on an ESX host, you can grant permissions to individual users and groups on that host. If you are acting in the Administrator role in vCenter Server, you can grant permissions to any user or group included in the Windows domain list that vCenter Server references.

vCenter Server 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 vCenter 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. Alternately, you can use the vSphere Client to delete the Windows Administrators group access permissions, but you must first create another permission at the root level that has a different user assigned to the Administrator role.

Working with Users and Groups on ESX Hosts

If you are directly connected to an ESX host through the vSphere Client, you can create, edit, and delete users and groups. These users and groups are visible in the vSphere Client whenever you log in to the ESX host, but are not available if you log in to vCenter Server.

View, Sort, and Export a List of Users and Groups

You can view, sort, and export lists of ESX users and groups to a file that is in HTML, XML, Microsoft Excel, or CSV format.

Procedure

- 1 Log in to the host using the vSphere Client.
- 2 Click the **Users & Groups** tab and click **Users** or **Groups**.

- 3 Determine how to sort the table, and hide or show columns according to the information you want to see in the exported file.
 - 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 show or hide columns, right-click any of the column headings and select or deselect the name of the column to hide.
- 4 Right-click anywhere in the table and click **Export List** to open the Save As dialog box.
- 5 Select a path and enter a filename.
- 6 Select the file type and click **OK**.

Add a User to the Users Table

Adding a user to the users table updates the internal user list that ESX maintains.

Procedure

- 1 Log in to the host using the vSphere Client.
- 2 Click the **Users & Groups** tab and click **Users**.
- 3 Right-click anywhere in the Users table and click **Add** to open the Add New User dialog box.
- 4 Enter a login, a user name, a numeric user ID (UID), and a password.

Specifying the user name and UID are optional. If you do not specify the UID, the vSphere Client assigns the next available UID.

Create a password that meets the length and complexity requirements. However, the ESX 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.

- 5 To allow a user to access the ESX host through a command shell, select **Grant shell access to this user**.
In general, do not grant shell access unless the user has a justifiable need. Users that access the host only through the vSphere Client do not need shell access.
- 6 To add the user to a group, select the group name from the **Group** drop-down menu and click **Add**.
- 7 Click **OK**.

Modify the Settings for a User

You can change the user ID, user name, password, and group settings for a user. You can also grant a user shell access.

Procedure

- 1 Log in to the host using the vSphere Client.
- 2 Click the **Users & Groups** tab and click **Users**.
- 3 Right-click the user and click **Edit** to open the Edit User dialog box.
- 4 To change the user ID, enter a numeric user UID in the **UID** text box.

The vSphere Client assigns the UID when you first create the user. In most cases, you do not have to change this assignment.

- 5 Enter a new user name.

- 6 To change the user's password, select **Change Password** and enter the new password.
- 7 To change the user's ability to access the ESX host through a command shell, select or deselect **Grant shell access to this user**.
- 8 To add the user to a group, select the group name from the **Group** drop-down menu and click **Add**.
- 9 To remove the user from a group, select the group name from the **Group membership** box and click **Remove**.
- 10 Click **OK**.

Remove a User or Group

You can remove a user or group from the ESX host.



CAUTION Do not remove the root user.

Procedure

- 1 Log in to the host using the vSphere Client.
- 2 Click the **Users & Groups** tab and click **Users** or **Groups**.
- 3 Right-click the user or group to remove and select **Remove**.

Add a Group to the Groups Table

Adding a group to the ESX groups table updates the internal group list maintained by the host.

Procedure

- 1 Log in to the host using the vSphere Client.
- 2 Click the **Users & Groups** tab and click **Groups**.
- 3 Right-click anywhere in the Groups table and click **Add** to open the Create New Group dialog box.
- 4 Enter a group name and numeric group ID (GID) in the **Group ID** text box.
Specifying the GID is optional. If you do not specify a GID, the vSphere Client assigns the next available group ID.
- 5 For each user that you want to add as a group member, select the user name from the list and click **Add**.
- 6 Click **OK**.

Add or Remove Users from a Group

You can add or remove a user from a group in the groups table.

Procedure

- 1 Log in to the host using the vSphere Client.
- 2 Click the **Users & Groups** tab and click **Groups**.
- 3 Right-click the group to modify and select **Properties** to open the Edit Group dialog box.
- 4 To add the user to a group, select the group name from the **Group** drop-down menu and click **Add**.
- 5 To remove the user from a group, select the group name from the **Group membership** box and click **Remove**.
- 6 Click **OK**.

Encryption and Security Certificates for ESX

ESX supports SSL v3 and TLS v1, generally referred to here as SSL. If SSL is enabled, data is private, protected, and 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.

Host certificate checking is enabled by default and SSL certificates are used to encrypt network traffic. However, ESX uses automatically generated certificates that are created as part of the installation process and stored on the host. These certificates are unique and make it possible to begin using the server, but they are not verifiable and are not signed by a trusted-well-known certificate authority (CA). These default certificates are vulnerable to possible man-in-the-middle attacks.

To receive the full benefit of certificate checking, particularly if you intend to use encrypted remote connections externally, install new certificates that are signed by a valid internal certificate authority or purchase a certificate from a trusted security authority.

NOTE 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. If CA-signed certificates are not installed, all communication between vCenter Server and vSphere Clients is encrypted using a self-signed certificate. These certificates do not provide the authentication security you might need in a production environment.

The default location for your certificate is `/etc/vmware/ssl/` on the ESX host. The certificate consists of two files: the certificate itself (`ru1.crt`) and the private-key file (`ru1.key`).

Enable Certificate Checking and Verify Host Thumbprints

To prevent man-in-the-middle attacks and to fully use the security that certificates provide, certificate checking is enabled by default. You can verify that certificate checking is enabled in the vSphere Client.

NOTE vCenter Server certificates are preserved across upgrades.

Procedure

- 1 Log in to a vCenter Server system using the vSphere Client.
- 2 Select **Administration > vCenter Server Settings**.
- 3 Click **SSL Settings** in the left pane and verify that **Check host certificates** is selected.
- 4 If there are hosts that require manual validation, compare the thumbprints listed for the hosts to the thumbprints in the host console.

To obtain the host thumbprint for ESX, run the following command.

```
openssl x509 -in /etc/vmware/ssl/ru1.crt -fingerprint -sha1 -noout
```

- 5 If the thumbprint matches, select the **Verify** check box next to the host.
Hosts that are not selected will be disconnected after you click **OK**.
- 6 Click **OK**.

Generate New Certificates for the ESX Host

The ESX host generates certificates the first time the system is started. Under certain circumstances, you might be required to force the host to generate new certificates. You typically generate new certificates only if you change the host name or accidentally delete the certificate.

Each time you restart the `vmware-hostd` process, the `mgmt-vmware` script searches for existing certificate files (`rui.crt` and `rui.key`). If it cannot find them, it generates new certificate files.

Procedure

- 1 Put the host into Maintenance Mode.
- 2 In the directory `/etc/vmware/ssl`, back up any existing certificates by renaming them using the following commands.

```
mv /etc/vmware/ssl/rui.crt /etc/vmware/ssl/orig.rui.crt
mv /etc/vmware/ssl/rui.key /etc/vmware/ssl/orig.rui.key
```

NOTE If you are regenerating certificates because you accidentally deleted them, you are not required to rename them.

- 3 Reboot your host to allow it to begin using the new certificate, or restart the host services:
 - a Use the following command to restart the `vmware-hostd` process.


```
service mgmt-vmware restart
```
 - b Use the following command to restart the `vmware-vmkauthd` process.


```
service vmware-vmkauthd restart
```
- 4 Confirm that the ESX host successfully generated new certificates by using the following command and comparing the time stamps of the new certificate files with `orig.rui.crt` and `orig.rui.key`.


```
ls -la /etc/vmware/ssl/rui*
```
- 5 Exit Maintenance Mode.

Replace a Default Certificate with a CA-Signed Certificate

The ESX host uses automatically generated certificates that are created as part of the installation process. These certificates are unique and make it possible to begin using the server, but they are not verifiable and they are not signed by a trusted, well-known certificate authority (CA). Using default certificates might not comply with the security policy of your organization. If you require a certificate from a trusted certificate authority, you can replace the default certificate.

Procedure

- 1 Put the host into Maintenance Mode.
- 2 Log in to the service console and acquire root privileges.
- 3 In the directory `/etc/vmware/ssl`, rename the existing certificates using the following commands.


```
mv rui.crt orig.rui.crt
mv rui.key orig.rui.key
```
- 4 Copy the new certificate and key to `/etc/vmware/ssl`.
- 5 Rename the new certificate and key to `rui.crt` and `rui.key`.

- 6 Reboot your host to allow it to begin using the new certificate, or restart the host services:
 - a Use the following command to restart the `vmware-hostd` process.


```
service mgmt-vmware restart
```
 - b Use the following command to restart the `vmware-vmkauthd` process.


```
service vmware-vmkauthd restart
```
- 7 Exit Maintenance Mode.

Configure SSL Timeouts

You can configure SSL timeouts for ESX.

Timeout periods can be set for two types of idle connections:

- The Read Timeout setting applies to connections that have completed the SSL handshake process with port 443 of ESX.
- The Handshake Timeout setting applies to connections that have not completed the SSL handshake process with port 443 of ESX.

Both connection timeouts are set in milliseconds.

Idle connections are disconnected after the timeout period. By default, fully established SSL connections have a timeout of infinity.

Procedure

- 1 Log in to the service console and acquire root privileges.
- 2 Change to the directory `/etc/vmware/hostd/`.
- 3 Use a text editor to open the `config.xml` file.
- 4 Enter the `<readTimeoutMs>` value in milliseconds.

For example, to set the Read Timeout to 20 seconds, enter the following command.

```
<readTimeoutMs>20000</readTimeoutMs>
```

- 5 Enter the `<handshakeTimeoutMs>` value in milliseconds.

For example, to set the Handshake Timeout to 20 seconds, enter the following command.

```
<handshakeTimeoutMs>20000</handshakeTimeoutMs>
```

- 6 Save your changes and close the file.
- 7 Enter the following command to restart the `vmware-hostd` process.

```
service mgmt-vmware restart
```

Example: Configuration File

The following section from the file `/etc/vmware/hostd/config.xml` shows where to enter the SSL timeout settings.

```
<vmacore>
...
<http>
    <readTimeoutMs>20000</readTimeoutMs>
</http>
...
<ssl>
```

```

...
<handshakeTimeoutMs>20000</handshakeTimeoutMs>
...
</ssl>
</vmacore>

```

Modifying ESX Web Proxy Settings

When you modify Web proxy settings, you have several encryption and user security guidelines to consider.

NOTE Restart the `vmware-hostd` process after making any changes to host directories or authentication mechanisms by entering the command `service mgmt-vmware restart`.

- Do not set up certificates using pass phrases. ESX does not support pass phrases, also known as encrypted keys. If you set up a pass phrase, ESX 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 If certificates are not stored locally on the host—for example, if they are stored on an NFS share—the host cannot access those certificates if ESX loses network connectivity. As a result, a client connecting to the host cannot successfully participate in a secure SSL handshake with the host.

- To support encryption for user names, passwords, and packets, SSL is enabled by default for vSphere Web Access and vSphere Web services SDK connections. To configure these connections so that they do not encrypt transmissions, disable SSL for your vSphere Web Access connection or vSphere Web Services SDK connection by switching the connection from HTTPS to HTTP.

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 services, such as the internal Web server that hosts vSphere Web Access, most internal ESX services are accessible only through port 443, the port used for HTTPS transmission. Port 443 acts as a reverse proxy for ESX. You can see a list of services on ESX through an HTTP welcome page, but you cannot directly access these services without proper authorization.

You can change this configuration so that individual services are directly accessible through HTTP connections. Do not make this change unless you are using ESX in a fully trusted environment.

- When you upgrade vCenter Server and vSphere Web Access, the certificate remains in place. If you remove vCenter Server and vSphere Web Access, the certificate directory is not removed from the service console.

Configure the Web Proxy to Search for Certificates in Nondefault Locations

You can configure the Web proxy so that it searches for certificates in a location other than the default location. This is useful for companies that centralize their certificates on a single machine so that multiple hosts can use the certificates.

Procedure

- 1 Log in to the service console and acquire root privileges.
- 2 Change to the `/etc/vmware/hostd/` directory.

- 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 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 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
```

Change Security Settings for a Web Proxy Service

You can change the security configuration so that individual services are directly accessible through HTTP connections.

Procedure

- 1 Log in to the service console and acquire root privileges.
- 2 Change to the `/etc/vmware/hostd/directory`.
- 3 Use a text editor to open the `proxy.xml` file.

The 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.
- `<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 are 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.
- `<serverNamespace>` is the namespace for the server that provides this service, for example `/sdk` or `/mob`.

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 vSphere Web Access to Communicate Through an Insecure Port

vSphere Web Access normally communicates with an ESX 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 Web server in `proxy.xml` file. In the following result, the access mode is changed from `httpsWithRedirect` to `httpAndHttps`.

```
<_type>vim.ProxyService.LocalServiceSpec</_type>  
<accessMode>httpAndHttps</accessMode>  
<port>8080</port>  
<serverNamespace>/ui</serverNamespace>
```

Service Console Security

VMware has basic security recommendations for using the service console, including how to use some of the service console's built-in security features. The service console is a management interface to ESX 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.

This chapter includes the following topics:

- [“General Security Recommendations,”](#) on page 175
- [“Log In to the Service Console,”](#) on page 176
- [“Service Console Firewall Configuration,”](#) on page 176
- [“Password Restrictions,”](#) on page 179
- [“Cipher Strength,”](#) on page 185
- [“setuid and setgid Flags,”](#) on page 185
- [“SSH Security,”](#) on page 187
- [“Security Patches and Security Vulnerability Scanning Software,”](#) on page 188

General Security Recommendations

To protect the service console against unauthorized intrusion and misuse, VMware imposes constraints on several service console parameters, settings, and activities. You can loosen the constraints to meet your configuration needs, but if you do so, make sure that 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 host.

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. Therefore, provide only trusted users with login access. By default, root access is limited by not allowing secure shell (SSH) login as the root user. Strongly consider keeping this default. Require ESX system administrators 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 vSphere Client to administer your ESX hosts.

Whenever possible, use vSphere Client, vSphere Web Access, or a third-party network management tool to administer your ESX hosts instead of working through the command-line interface as the root user. Using vSphere 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 do not need.

- Use only VMware sources to upgrade ESX 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 must 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.

Log In to the Service Console

Although you perform most ESX configuration activities through the vSphere 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.

Procedure

- 1 Log in to the ESX host using one of the following methods.
 - If you have direct access to the host, press Alt+F2 to open the login page on the machine's physical console.
 - If you are connecting to the host remotely, use SSH or another remote console connection to start a session on the host.
- 2 Enter a user name and password recognized by the ESX host.

If you are performing activities that require root privileges, log in to the service console as a recognized user and acquire root privileges through the `sudo` command, which provides enhanced security compared to the `su` command.

What to do next

In addition to ESX-specific commands, you can use the service console command-line interface to run many Linux and UNIX commands. For more information about service console commands, use the `man <command_name>` command to check for man pages.

Service Console Firewall Configuration

ESX 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 ports 22, 123, 427, 443, 902, 5989, 5988, which are used for basic communication with ESX. This setting enforces a high level of security for the 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 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 be required to open additional ports after installation. For a list of commonly used ports that you might open, see [“TCP and UDP Ports for Management Access,”](#) on page 149.

As you add the supported services and management agents required to operate ESX effectively, you open other ports in the service console firewall. You add services and management agents through vCenter Server as described in [“Configuring Firewall Ports for Supported Services and Management Agents,”](#) on page 146.

In addition to the ports you open for these services and agents, you might 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.

Determine the Service Console Firewall 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.

Procedure

- 1 Log in to the service console and acquire root privileges.
- 2 Use the following two commands to determine whether incoming and outgoing traffic is blocked or allowed.

```
esxcfg-firewall -q incoming
esxcfg-firewall -q outgoing
```

Interpret the results according to [Table 14-1](#).

Table 14-1. Service Console Firewall 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

Set the Service Console Firewall Security Level

After you determine the level of firewall security for the service console, you can set the security level. 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.

Procedure

- 1 Log in to the service console and acquire root privileges.

- 2 Run one of the following commands to set the service console firewall security level.

- 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 Use 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 did not explicitly open, raising the firewall setting to high does not terminate the backup. The backup completes, releases the connection, and no further connections are accepted for the port.

Open a Port 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.

Prerequisites

Use this procedure only to open ports for services or agents that are not configurable through the vSphere Client.



CAUTION VMware supports opening and closing firewall ports only through the vSphere Client or the `esxcfg-firewall` command. Using any other methods or scripts to open firewall ports can lead to unexpected behavior.

Procedure

- 1 Log in to the service console and acquire root privileges.

- 2 Use the following command to open the port.

```
esxcfg-firewall --openPort <port_number>,tcp|udp,in|out,<port_name>
```

- `<port_number>` is the vendor-specified port number.
- Use `tcp` for TCP traffic or `udp` for UDP traffic.
- Use `in` to open the port for inbound traffic or `out` to open it for outbound traffic.
- `<port_name>` is a descriptive name to help identify the service or agent using the port. A unique name is not required.

For example:

```
esxcfg-firewall --openPort 6380,tcp,in,Navisphere
```

- 3 Run the following command to restart the `vmware-hostd` process.

```
service mgmt-vmware restart
```

Close a Port in the Service Console Firewall

You can close particular ports in the service console firewall. If you close a port, active sessions of the service associated with the port are not necessarily 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.

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 vSphere Client.

Prerequisites

Use this procedure only to close ports for services or agents not specifically configurable through the vSphere Client.



CAUTION VMware supports opening and closing firewall ports only through the vSphere Client or the `esxcfg-firewall` command. Using any other methods or scripts to open and close firewall ports can lead to unexpected behavior.

Procedure

1 Log in to the service console and acquire root privileges.

2 Use the following command to close the port.

```
esxcfg-firewall --closePort <port_number>,tcp|udp,in|out,<port_name>
```

The `<port_name>` argument is optional.

For example:

```
esxcfg-firewall --closePort 6380,tcp,in
```

3 Use the following command to restart the `vmware-hostd` process.

```
service mgmt-vmware restart
```

Password Restrictions

The ease with which an attacker can log in to an ESX host depends on finding a legitimate user name and password combination. You can set password restrictions to help prevent attackers from obtaining user passwords.

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 to try every character combination up to a certain length or use 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 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 repeatedly.

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, which eliminates any benefit.

To help protect your password database from misuse, password shadowing is enabled so that password hashes are hidden from access. Also, ESX uses MD5 password hashes, which provide stronger password security and lets you set minimum length requirements to more than eight characters.

Password Aging

You can impose password aging restrictions to ensure that user passwords do not stay active for long periods.

ESX imposes the following password aging restrictions for user logins by default.

Maximum days	The number of days that a user can keep a password. By default, passwords are set to never expire.
Minimum days	The minimum number of days between password changes. The default is 0, meaning that the users can change their passwords any time.
Warning time	The number of days in advance of password expiration that a reminder is sent. The default is seven days. Warnings are only displayed when logging directly in to the service console or when using SSH.

You can tighten or loosen any of these settings. You can also override the default password aging settings for an individual user or group.

Change Default Password Aging Restrictions for a Host

You can impose stricter or looser password aging restrictions for a host than those provided by default.

Procedure

- 1 Log in to the service console and acquire root privileges.
- 2 To change the maximum number of days a user can keep a password, use the following command.
`esxcfg-auth --passmaxdays=<number_of_days>`
- 3 To change the minimum number of days between password changes, use the following command.
`esxcfg-auth --passmindays=<number_of_days>`
- 4 To change the warning time before a password change, use the following command.
`esxcfg-auth --passwarnage=<number_of_days>`

Change Default Password Aging Restrictions for Users

You can override the default password aging restrictions for particular users or groups.

Procedure

- 1 Log in to the service console and acquire root privileges.
- 2 To change the maximum number of days, use the following command.
`chage -M <number_of_days> <username>`
- 3 To change the warning time, use the following command.
`chage -W <number_of_days> <username>`
- 4 To change the minimum number of days, use the following command.
`chage -m <number_of_days> <username>`

Password Complexity

By default, ESX 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 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. ESX has defaults for password standards and retry restrictions.

Minimum length

The minimum password length 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 14-2 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, so although the mathematically accurate character requirement for a four character-class password is five characters, the effective requirement is six.

Table 14-2. 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 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 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 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.

- 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.

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. You cannot specify these additional parameters in `esxcfg-auth`.

Configure a Password Reuse Rule

You can set the number of old passwords that are stored for each user.

Procedure

- 1 Log in to the service console and acquire root privileges.
- 2 Change to the directory `/etc/pam.d/`.
- 3 Use a text editor to open the `system-auth-generic` file.
- 4 Locate the line that starts with `password sufficient /lib/security/$ISA/pam_unix.so`.
- 5 Add the following parameter to the end of the line, where X is the number of old passwords to store for each user.

`remember=X`

 Use a space between parameters.
- 6 Save your changes and close the file.
- 7 Change to the directory `/etc/security/` and use 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
```

Change Default Password Complexity for the pam_cracklib.so Plug-In

The `pam_cracklib.so` plug-in allows you to set the minimum length and complexity of a password.

To make a password more complex, you can assign values to the credit parameters for each of the following character classes:

- `<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 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 require 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>`.

Procedure

- 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>
<retries> is the number of retries users are allowed before they are locked out.
```

Example: 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:

$$(x + y + z + p + q + e + \#) + (lc_credit + oc_credit) = 9$$

While the password score is 9, it does not contain the required uppercase letter and digit.

The password candidate **Xyzpq3#** would be accepted:

$$(X + y + z + p + q + 3 + \#) + (lc_credit + oc_credit) = 9$$

The password score for this example is also 9, but this password includes the required uppercase letter and digit. The uppercase letter and digit do not add extra credit.

Switch to the pam_passwdqc.so Plug-In

The `pam_cracklib.so` plug-in provides sufficient password strength enforcement for most environments. However, if the plug-in is not stringent enough for your needs, you can use the `pam_passwdqc.so` plug-in instead.

The `pam_passwdqc.so` 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.

NOTE The `pam_passwdqc.so` plug-in used in Linux provides more parameters than the parameters supported for ESX. You cannot specify these additional parameters in `esxcfg-auth`. For more information on this plug-in, see your Linux documentation.

Procedure

- 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>
```

- `<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 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 use 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 cannot 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 ciphers are 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 uses one of the strongest block ciphers available—256-bit AES block encryption. ESX also uses 1024-bit RSA for key exchange. These encryption algorithms are the default for the following connections.

- vSphere Client connections to vCenter Server and to the ESX host through the service console.
- vSphere Web Access connections to the ESX host through the service console.

NOTE Because use of vSphere Web Access ciphers is determined by the Web browser you are using, this management tool might use other ciphers.

- SDK connections to vCenter Server and to ESX.
- Service console connections to virtual machines through the VMkernel.
- SSH connections to the ESX host through the service console.

setuid and setgid Flags

During ESX installation, several applications that include the `setuid` and `setgid` flags are installed by default. Some of the applications provide facilities required for correct operation of the host. Others are optional, but they can make maintaining and troubleshooting the host and the network easier.

setuid	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	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.

Disable Optional Applications

Disabling any of the required applications results in problems with ESX authentication and virtual machine operation, but you can disable any optional application.

Optional applications are listed in [Table 14-3](#) and [Table 14-4](#).

Procedure

- 1 Log in to the service console and acquire root privileges.
- 2 Run one of the following commands to disable the application.
 - For setuid flagged applications:
`chmod a-s <path_to_executable_file>`
 - For setgid flagged applications:
`chmod a-g <path_to_executable_file>`

Default setuid Applications

Several applications that include the setuid flag are installed by default.

[Table 14-3](#) lists the default setuid applications and indicates whether the application is required or optional.

Table 14-3. 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

Table 14-3. Default setuid Applications (Continued)

Application	Purpose and Path	Required or Optional
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
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

Default setgid Applications

Two applications that include the setgid flag are installed by default.

[Table 14-4](#) lists the default setgid applications and indicates whether the application is required or optional.

Table 14-4. 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

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 the host. SSH is used for secure logins and data transfers because it offers stronger protection than other command shells.

In this ESX release, the SSH configuration is enhanced to provide a higher security level. This enhancement includes the following key features.

- 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 run 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.

Change the Default SSH Configuration

You can change the default SSH configuration.

Procedure

- 1 Log in to the service console and acquire root privileges.
- 2 Change to the `/etc/ssh` directory.
- 3 Use a text editor to perform any of the following actions in the `sshd_config` file.
 - To allow remote root login, change the setting to `yes` in the following line.
`PermitRootLogin no`
 - To revert to the default SSH protocol (Version 1 and 2), comment out the following line.
`Protocol 2`
 - To revert to the 3DES cipher and other ciphers, comment out the following line.
`Ciphers aes256-cbc,aes128-cbc`
 - To disable Secure FTP (SFTP) on SSH, comment out the following line.
`Subsystem ftp /usr/libexec/openssh/sftp-server`
- 4 Save your changes and close the file.
- 5 Run the following command to restart the SSHD service.
`service sshd restart`

Security Patches and Security Vulnerability Scanning Software

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 does not include the most recent security patches even though it does. If this occurs, you can perform certain checks.

This problem is common to the industry and not specific to VMware. 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 does not report these false positives.

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 a bulletin that contains a list of vSphere Installation Bundles (VIBs) that you use to update the software on ESX. Although these fixes might be available from other sources, always use bulletins that VMware generates instead of using third-party RPM Package Manager packages.

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.

The following is an example of how this problem occurs:

- 1 You initially install ESX 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 a VIB. The OpenSSL version in the VIB is 0.9.7a-1, indicating that the original version (0.9.7a) now contains patch 1.
- 4 You install the updates.
- 5 The security scanner fails to note the -1 suffix and erroneously reports that security for OpenSSL is not 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 if you require an update.
- Read the VMware VIB documentation for information on the patch contents.
- Look for the Common Vulnerabilities and Exposures (CVE) number from the security alert in the software update change log.

If the CVE number is there, the specified package addresses that vulnerability.

Security Deployments and Recommendations

15

ESX deployment scenarios can help you understand how best to employ the security features in your own deployment. Scenarios also illustrate some basic security recommendations that you can consider when creating and configuring virtual machines.

This chapter includes the following topics:

- [“Security Approaches for Common ESX Deployments,”](#) on page 191
- [“Virtual Machine Recommendations,”](#) on page 195

Security Approaches for Common ESX Deployments

You can compare security approaches for different types of deployments to help plan security for your own ESX deployment.

The complexity of ESX deployments can vary significantly depending on the size of your company, the way that data and resources are shared with the outside world, whether there are multiple datacenters or only one, and so forth. Inherent in the following deployments are policies for user access, resource sharing, and security level.

Single-Customer Deployment

In a single-customer deployment, ESX hosts are owned and maintained within a single corporation and single datacenter. Host resources are not shared with outside users. One site administrator maintains the hosts, which are run on a number of virtual machines.

The single-customer 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 host and cannot access any of the ESX tools such as vCenter Server 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 15-1](#) shows how you might handle sharing for the components that you use and configure for the host.

Table 15-1. 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 must use this VLAN.

Table 15-1. Sharing for Components in a Single-Customer Deployment (Continued)

Function	Configuration	Comments
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 must use this switch or adapter. You can configure your virtual machines on the same virtual switch and network adapter.
VMFS sharing?	Yes	All .vmdk files 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 176 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 15-2](#) shows how you might set up user accounts for the host.

Table 15-2. 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 15-3](#) shows the level of access for each user.

Table 15-3. User Access in a Single-Customer Deployment

Access Level	Site Administrator	System Administrator
Root access?	Yes	No
Service console access through SSH?	Yes	No
vCenter Server and vSphere Web Access?	Yes	No
Virtual machine creation and modification?	Yes	No
Virtual machine access through the console?	Yes	Yes

Multiple-Customer Restricted Deployment

In a multiple-customer restricted deployment, ESX hosts are in the same datacenter and are used to serve applications for multiple customers. The site administrator maintains the 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 host, but the site administrator restricts resource sharing to prevent rogue interaction.

Although 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 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 15-4](#) shows how you might handle sharing for the components you use and configure for the host.

Table 15-4. 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 must 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 must use this switch or adapter. You configure virtual machines for one customer so that they all share the same virtual switch and network adapter. They do not share the switch and adapter with any other customers.
VMFS sharing?	No	Each customer has its own VMFS partition, and the 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 15-5 shows how you might set up user accounts for the ESX host.

Table 15-5. 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 15-6 shows the level of access for each user.

Table 15-6. 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
vCenter Server and vSphere 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 a multiple-customer open deployment, ESX hosts are in the same datacenter and are used to serve applications for multiple customers. The site administrator maintains the 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 host, but there are fewer restrictions on resource sharing.

Although there is only one site administrator in a multiple-customer open deployment, several customer administrators maintain the virtual machines assigned to their customers. The deployment also includes customer system administrators who do not have ESX 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 15-7](#) shows how you might handle sharing for the components that you use and configure for the host.

Table 15-7. 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 must 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 must 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 15-8](#) shows how you might set up user accounts for the host.

Table 15-8. 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 15-9](#) shows the level of access for each user.

Table 15-9. 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
vCenter Server and vSphere 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

There are several safety precautions to consider when evaluating virtual machine security and administering virtual machines.

Installing Antivirus Software

Because each virtual machine hosts a standard operating system, 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.

Stagger the schedule for virus scans, particularly in deployments with a large number of virtual machines. Performance of systems in your environment will degrade significantly if you scan all virtual machines simultaneously.

Because software firewalls and antivirus software can be virtualization-intensive, you can balance the need for these two security measures against virtual machine performance, especially if you are confident that your virtual machines are in a fully trusted environment.

Disable Copy and Paste Operations Between the Guest Operating System and Remote Console

You can disable copy and paste operations to prevent exposing sensitive data that has been copied to the clipboard.

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.

Procedure

- 1 Log in to a vCenter Server system using the vSphere Client.
- 2 On the **Summary** tab, click **Edit Settings**.
- 3 Select **Options > Advanced > General** and click **Configuration Parameters**.

- Click **Add Row** and type the following values in the Name and Value columns.

Name	Value
isolation.tools.copy.disable	true
isolation.tools.paste.disable	true
isolation.tools.setGUIOptions.enable	false

NOTE These options override any settings made in the guest operating system's VMware Tools control panel.

These selections create the following field values.

Name	Field Value
sched.mem.max	unlimited
sched.swap.derivedName	/vmfs/volumes/e5f9f3d1-ed4d8ba/New Virtual Machine
scsi0:0.redo	true
vmware.tools.installstate	none
vmware.tools.lastInstallStatus.result	unknown
isolation.tools.copy.disable	true
isolation.tools.paste.disable	true
isolation.tools.setGUIOptions.enable	false

- Click **OK** to close the Configuration Parameters dialog box, and click **OK** again to close the Virtual Machine Properties dialog box.

Removing Unnecessary Hardware Devices

Users and processes without privileges on a virtual machine can connect or disconnect hardware devices, such as network adapters and CD-ROM drives. Therefore, removing unnecessary hardware devices can help prevent attacks.

Attackers can use this capability to breach virtual machine security in several ways. For example, 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, or 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 vSphere Client **Configuration** tab to remove any unneeded or unused hardware devices. Although this measure tightens virtual machine security, it is not a good solution in situations where you might bring an unused device back into service at a later time.

Prevent a Virtual Machine User or Process from Disconnecting Devices

If you do not 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.

Procedure

- Log in to a vCenter Server system using the vSphere Client.
- Select the virtual machine in the inventory panel.
- On the **Summary** tab, click **Edit Settings**.

- 4 Select **Options > General Options** and make a record of the path displayed in the **Virtual Machine Configuration File** text box.

- 5 Log in to the service console and acquire root privileges.

- 6 Change directories to access the virtual machine configuration file whose path you recorded in [Step 4](#).

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 would change to the following directory.

```
/vmfs/volumes/vol1/vm-finance/
```

- 7 Use a text editor to add the following line to the `.vmx` file, where `<device_name>` is the name of the device you want to protect (for example, `ethernet1`).

```
<device_name>.allowGuestConnectionControl = "false"
```

NOTE By default, Ethernet 0 is configured to disallow device disconnection. The only reason you might change this is if a prior administrator set `<device_name>.allowGuestConnectionControl` to `true`.

- 8 Save your changes and close the file.

- 9 In the vSphere Client, right-click the virtual machine and select **Power Off**.

- 10 Right-click the virtual machine and select **Power On**.

Limiting Guest Operating System Writes to Host Memory

The guest operating system processes send informational messages to the ESX host through VMware Tools. 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 denial-of-service (DoS) attack.

The informational messages sent by guest operating processes are known as `setinfo` messages and typically contain name-value pairs that define virtual machine characteristics or identifiers that the host stores—for example, `ipaddress=10.17.87.224`. The configuration file containing these name-value pairs is limited to a size of 1MB, which prevents attackers from staging a DoS attack by writing software that mimics VMware Tools and filling the host's memory with arbitrary configuration data, which consumes space needed by the virtual machines.

If you require more than 1MB of storage for name-value pairs, you can change the value as required. You can also prevent the guest operating system processes from writing any name-value pairs to the configuration file.

Modify Guest Operating System Variable Memory Limit

You can increase the guest operating system variable memory limit if large amounts of custom information are being stored in the configuration file.

Procedure

- 1 Log in to a vCenter Server system using the vSphere Client.
- 2 Select the virtual machine in the inventory panel.
- 3 On the **Summary** tab, click **Edit Settings**.
- 4 Select **Options > Advanced > General** and click **Configuration Parameters**.

- 5 If the size limit attribute is not present, you must add it.
 - a Click **Add Row**.
 - b In the Name column, type **tools.setInfo.sizeLimit**.
 - c In the Value column, type **Number of Bytes**.

If the size limit attribute exists, modify it to reflect the appropriate limits.

- 6 Click **OK** to close the Configuration Parameters dialog box, and click **OK** again to close the Virtual Machine Properties dialog box.

Prevent the Guest Operating System Processes from Sending Configuration Messages to the Host

You can 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.

Procedure

- 1 Log in to a vCenter Server system using the vSphere Client.
- 2 Select the virtual machine in the inventory panel.
- 3 On the **Summary** tab, click **Edit Settings**.
- 4 Select **Options > Advanced > General** and click **Configuration Parameters**.
- 5 Click **Add Row** and type the following values in the Name and Value columns.
 - In the Name column: **isolation.tools.setinfo.disable**
 - In the Value column: **true**
- 6 Click **OK** to close the Configuration Parameters dialog box, and 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 you reboot a host, so the file can grow to be quite large. You can ensure new log file creation happens more frequently by limiting the maximum size of the log files. VMware recommends saving 10 log files, each one limited to 100KB. These values are large enough to 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. If it is over the limit, the next entry is written to a new log. If the maximum number of log files exists, the oldest log file is deleted. A DoS attack that avoids these limits could be attempted by writing an enormous log entry, but each log entry is limited in size to 4KB, so no log files are ever more than 4KB larger than the configured limit.

Limit Log File Numbers and Sizes

To prevent virtual machine users and processes from flooding the log file, which can lead to denial of service, you can limit the number and size of the log files ESX generates.

Procedure

- 1 Log in to a vCenter Server system using the vSphere Client.
- 2 On the **Summary** tab, click **Edit Settings**.
- 3 Select **Options > General Options** and make a record of the path displayed in the **Virtual Machine Configuration File** text box.
- 4 Log into 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 would change to the following directory.

```
/vmfs/volumes/vol1/vm-finance/
```

- 6 To limit the log size, use a text editor to add or edit the following line to the `.vmx` file, where `<maximum_size>` is the maximum file size in bytes.

```
log.rotateSize=<maximum_size>
```

For example, to limit the size to around 100KB, enter **100000**.

- 7 To keep a limited number of log files, use a text editor to add or edit the following line to the `.vmx` file, where `<number_of_files_to_keep>` is the number of files the server keeps.

```
log.keepOld=<number_of_files_to_keep>
```

For example, to keep 10 log files and begin deleting the oldest ones as new ones are created, enter **10**.

- 8 Save your changes and close the file.

Disable Logging for the Guest Operating System

If you choose not to write troubleshooting information into a virtual machine log file stored on the VMFS volume, you can stop logging altogether.

If you disable logging for the guest operating system, 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.

Procedure

- 1 Log in to a vCenter Server system using the vSphere Client and select the virtual machine in the inventory.
- 2 On the **Summary** tab, click **Edit Settings**.
- 3 Click the **Options** tab and in the options list under Advanced, select **General**.
- 4 In Settings, deselect **Enable logging**.
- 5 Click **OK** to close the Virtual Machine Properties dialog box.

Host Profiles

Managing Host Profiles

The host profiles feature creates a profile that encapsulates the host configuration and helps to manage the host configuration, especially in environments where an administrator manages more than one host or cluster in vCenter Server.

Host profiles eliminates per-host, manual, or UI-based host configuration and maintain configuration consistency and correctness across the datacenter by using host profile policies. These policies capture the blueprint of a known, validated reference host configuration and use this to configure networking, storage, security, and other settings on multiple hosts or clusters. You can then check a host or cluster against a profile's configuration for any deviations.

This chapter includes the following topics:

- [“Host Profiles Usage Model,”](#) on page 203
- [“Access Host Profiles View,”](#) on page 204
- [“Creating a Host Profile,”](#) on page 204
- [“Export a Host Profile,”](#) on page 205
- [“Import a Host Profile,”](#) on page 205
- [“Edit a Host Profile,”](#) on page 206
- [“Manage Profiles,”](#) on page 207
- [“Checking Compliance,”](#) on page 210

Host Profiles Usage Model

This topic describes the workflow of using Host Profiles.

You must have an existing vSphere installation with at least one properly configured host.

- 1 Set up and configure the host that will be used as the reference host.
A reference host is the host from which the profile is created.
- 2 Create a profile using the designated reference host.
- 3 Attach a host or cluster with the profile.
- 4 Check the host's compliance against a profile. This ensures that the host continues to be correctly configured.

- 5 Apply the host profile of the reference host to other hosts or clusters of hosts.

NOTE Host profiles is only supported for VMware vSphere 4.0 hosts. This feature is not supported for VI 3.5 or earlier hosts. If you have VI 3.5 or earlier hosts managed by your vCenter Server 4.0, the following can occur if you try to use host profiles for those hosts:

- You cannot create a host profile that uses a VMware Infrastructure 3.5 or earlier host as a reference host.
- You cannot apply a host profile to any VI 3.5 or earlier hosts. The compliance check fails.
- While you can attach a host profile to a mixed cluster that contains VI 3.5 or earlier hosts, the compliance check for those hosts fails.

As a licensed feature of vSphere, Host Profiles are only available when the appropriate licensing is in place. If you see errors, please ensure that you have the appropriate vSphere licensing for your hosts.

Access Host Profiles View

The Host Profiles main view lists all available profiles. Administrators can also use the Host Profiles main view to perform operations on host profiles and configure profiles.

The Host Profiles main view should be used by experienced administrators who wish to perform host profile operations and configure advanced options and policies. Most operations such as creating new profiles, attaching entities, and applying profiles can be performed from the Hosts and Clusters view.

Procedure

- ◆ Select **View > Management > Host Profiles**.

Any existing profiles are listed on the left side in the profiles list. When a profile is selected from the profile list, the details of that profile are displayed on the right side.

Creating a Host Profile

You create a new host profile by using the designated reference host's configuration.

A host profile can be created from the Host Profiles main view or the host's context menu in the Hosts and Clusters.

Create a Host Profile from Host Profiles View

You can create a host profile from the Host Profiles main view by using the configuration of an existing host.

Prerequisites

You must have a vSphere installation and at least one properly configured host in the inventory.

Procedure

- 1 In the Host Profiles main view, click **Create Profile**.
The Create Profile wizard appears.
- 2 Select the option to create a new profile and click **Next**.
- 3 Select the host to use to create the profile and click **Next**.
- 4 Type the name and enter a description for the new profile and click **Next**.
- 5 Review the summary information for the new profile and click **Finish** to complete creating the profile.

The new profile appears in the profile list.

Create a Host Profile from Host

You can create a new host profile from the host's context menu in the Hosts and Clusters inventory view.

Prerequisites

You must have a vSphere installation and at least one properly configured host in the inventory.

Procedure

- 1 In the Hosts and Clusters view, select the host that you want to designate as the reference host for the new host profile.
- 2 Right-click the host and select **Host Profile > Create Profile from Host**.
The Create Profile from Host wizard opens.
- 3 Type the name and enter a description for the new profile and click **Next**.
- 4 Review the summary information for the new profile and click **Finish** to complete creating the profile.

The new profile appears in the host's Summary tab.

Export a Host Profile

You can export a profile to a file that is in the VMware profile format (.vpf).

Procedure

- 1 In the Host Profiles main page, select the profile to export from the profile list.
- 2 Right-click the profile and select **Export Profile**.
- 3 Select the location and type the name of the file to export the profile.
- 4 Click **Save**.

Import a Host Profile

You can import a profile from a file in the VMware profile format (.vpf).

Procedure

- 1 In the Host Profiles main page, click the **Create Profile** icon.
The Create Profile wizard appears.
- 2 Select the option to import a profile and click **Next**.
- 3 Enter or browse for the VMware Profile Format file to import and click **Next**.
- 4 Type the name and enter a description for the imported profile and click **Next**.
- 5 Review the summary information for the imported profile and click **Finish** to complete importing the profile.

The imported profile appears in the profile list.

Edit a Host Profile

You can view and edit host profile policies, select a policy to be checked for compliance, and change the policy name or description.

Procedure

- 1 In the Host Profiles main view, select the profile to edit from the profile list.
- 2 Click **Edit Host Profile**.
- 3 Change the profile name or description in the fields at the top of the Profile Editor.
- 4 Edit or disable the policy.
- 5 Enable the policy compliance check.
- 6 Click **OK** to close the Profile Editor.

Edit a Policy

A policy describes how a specific configuration setting should be applied. The Profile Editor allows you to edit policies belonging to a specific host profile.

On the left side of the Profile Editor, you can expand the host profile. Each host profile is composed of several sub-profiles that are designated by functional group to represent configuration instances. Each sub-profile contains many policies that describe the configuration that is relevant to the profile.

The sub-profiles (and example policies and compliance checks) that may be configured are:

Table 16-1. Host Profile Sub-Profile Configurations

Sub-Profile Configuration	Example Policies and Compliance Checks
Memory reservation	Set memory reservation to a fixed value.
Storage	Configure NFS storage.
Networking	Configure virtual switch, port groups, physical NIC speed, security and NIC teaming policies.
Date and Time	Configure time settings, timezone of server.
Firewall	Enable or disable a ruleset.
Security	Add a user or usergroup.
Service	Configure settings for a service.
Advanced	Modify advanced options.

Procedure

- 1 Open the Profile Editor for the profile you wish to edit.
- 2 On the left side of the Profile Editor, expand a sub-profile until you reach the policy you want to edit.
- 3 Select the policy.

On the right side of the Profile Editor, the policy options and parameters are displayed within the **Configuration Details** tab.

- 4 Select a policy option from the drop-down menu and set its parameter.
- 5 (Optional) If you make a change to a policy, but wish to revert back to the default option, click **Revert** and the option is reset.

Enable Compliance Check

You can decide whether a host profile policy is checked for compliance.

Procedure

- 1 Open the Profile Editor for a profile and navigate to the policy you wish to enable for compliance check.
- 2 On the right-hand side of the Profile Editor, select the **Compliance Details** tab.
- 3 Enable the check box for the policy.

NOTE If you disable the check box so this policy is not checked for compliance, the other policies that are enabled for compliance check will still be checked.

Manage Profiles

After you create a host profile, you can manage the profile by attaching a profile to a particular host or cluster and then applying that profile to the host or cluster.

Attaching Entities

Hosts that need to be configured are attached to a profile. Profiles can also be attached to a cluster. In order to be compliant, all hosts within an attached cluster must be configured according to the profile.

You can attach a host or cluster to a profile from:

- Host Profiles main view
- Host's context menu
- Cluster's context menu
- Cluster's Profile Compliance tab

Attach Entities from the Host Profiles View

Before you can apply the profile to an entity (host or cluster of hosts), you must attach the entity to the profile.

You can attach a host or cluster to a profile from the Host Profiles main view.

Procedure

- 1 In the Host Profiles main view, select the profile to which you want to add the attachment from the profile list.
- 2 Click the **Attach Host/Cluster** icon.
- 3 Select the host or cluster from the expanded list and click **Attach**.
The host or cluster is added to the Attached Entities list.
- 4 (Optional) Click **Detach** to remove an attachment from a host or cluster.
- 5 Click **OK** to close the dialog.

Attach Entities from the Host

Before you can apply the profile to host you must attach the entity to the profile.

You can attach a profile to a host from the host's context menu in the Hosts and Clusters inventory view.

Procedure

- 1 In the Hosts and Clusters view, select the host to which you want to attach a profile.
- 2 Right-click the host and select **Host Profile > Manage Profile**.

NOTE If no host profiles exist in your inventory, a dialog appears asking if you want to create and attach the host to this profile.

- 3 In the Change Attached Profile dialog, select the profile to attach to the host and click **OK**.

The host profile is updated in the **Summary** tab of the host.

Applying Profiles

To bring a host to the desired state as specified in the profile, apply the profile to the host.

You can apply a profile to a host from:

- Host Profiles main view
- Host's context menu
- Cluster's Profile Compliance tab

Apply a Profile from the Host Profiles View

You can apply a profile to a host from the Host Profiles main view.

Prerequisites

The host must be in maintenance mode before a profile is applied to it.

Procedure

- 1 In the Host Profiles main view, select the profile you want to apply to the host .
- 2 Select the **Hosts and Clusters** tab.

The list of attached hosts are shown under Entity Name.

- 3 Click **Apply Profile**.

In the Profile Editor, you might be prompted to enter the required parameters needed to apply the profile.

- 4 Enter the parameters and click **Next**.
- 5 Continue until all the required parameters are entered.
- 6 Click **Finish**.

Compliance Status is updated.

Apply a Profile from the Host

You can apply a profile to a host from the host's context menu.

Prerequisites

The host must be in maintenance mode before applying it to a profile.

Procedure

- 1 In the Hosts and Clusters view, select the host to which you want to apply a profile.
- 2 Right-click the host and select **Host Profile > Apply Profile**.

- 3 In the Profile Editor, enter the parameters and click **Next**.
- 4 Continue until all the required parameters are entered.
- 5 Click **Finish**.

Compliance Status is updated.

Change Reference Host

You can change the reference host in the host profile.

Procedure

- 1 You can perform this task either from the Host Profiles main view or from the host.
 - ◆ In the Host Profiles main view, right-click the profile you want to change and select **Change Reference Host**.
 - ◆ In the Hosts and Clusters view, right-click the host you want to change and select **Manage Profiles**.
- 2 In the Detach or Change Host Profile dialog box, select whether to detach the profile from the host or cluster or change the profile's reference host.
 - ◆ Click **Detach** to remove the association between the host and the profile.
 - ◆ Click **Change** to continue with updating the profile's reference host.

If you selected **Change**, the Change Reference Host dialog box opens. The current host that the profile references is displayed as the reference host.

- 3 Expand the inventory list and select the host to which to attach the profile.
- 4 Click **Update**.

The reference host is updated.

- 5 Click **OK**.

The Summary tab for the host profile lists the updated reference host.

Manage Profiles from a Cluster

You can create a profile, attach a profile, or update reference hosts from the cluster's context menu.

Procedure

- ◆ In the Hosts and Clusters view, right-click a cluster and select **Host Profile > Manage Profile**.

Depending on your host profile setup, perform one of the following actions.

Profile Status	Result
The cluster is not attached to a host profile and no profile exists in your inventory.	You are asked whether you want to create a profile. If you select Yes , the Create Profile wizard opens.
The cluster is not attached to a host profile and one or more profiles exist in your inventory.	The Attach Profile dialog box opens. Select the profile that you want to attach and click OK .
The cluster is already attached to a host profile.	Click Detach or click Change to attach a different profile to the cluster.

Checking Compliance

Checking compliance ensures that the host or cluster continues to be correctly configured.

After a host or cluster is configured with the reference host profile, a manual change, for example, can occur, making the configuration incorrect. Checking compliance on a regular basis ensures that the host or cluster continues to be correctly configured.

Check Compliance from the Host Profiles View

You can check the compliance of a host or cluster to a profile from the Host Profiles main view.

Procedure

- 1 From the Host Profiles list, select the profile that you want to check.
- 2 In the **Hosts and Clusters** tab, select the host or cluster from the list under Entity Name.
- 3 Click **Check Compliance Now**.

The compliance status is updated as Compliant, Unknown, or Non-compliant.

If the compliance status is Non-compliant, you can apply the host to the profile.

Check Compliance from Host

After a profile has been attached to a host, run a compliance check to verify the configuration.

Procedure

- 1 In the Hosts and Clusters view, select the host on which you want to run the compliance check.
- 2 Right-click the host and select **Host Profile > Check Compliance**.

The host's compliance status is displayed in the host's **Summary** tab.

If the host is not compliant, you must apply the profile to the host.

Check Cluster Compliance

A cluster may be checked for compliance against a host profile or for specific cluster requirements and settings.

Procedure

- 1 In the Hosts and Clusters view, select the cluster on which you want to run the compliance check.
- 2 In the Profile Compliance tab, click **Check Compliance Now** to check the cluster's compliance with both the host profile that is attached to this cluster and the cluster requirements, if any.
 - The cluster is checked for compliance with specific settings for hosts in the cluster, such as DRS, HA, and DPM. For example, it may check if VMotion is enabled. The compliance status for the cluster requirements is updated. This check is performed even if a host profile is not attached to the cluster.
 - If a host profile is attached to the cluster, the cluster is checked for compliance with the host profile. The compliance status for the host profile is updated.
- 3 (Optional) Click **Description** next to the Cluster Requirements for a list of the specific cluster requirements.
- 4 (Optional) Click **Description** next to Host Profiles for a list of the specific host profile compliance checks.
- 5 (Optional) Click **Change** to change the host profile that is attached to the cluster.
- 6 (Optional) Click **Remove** to detach the host profile that is attached to the cluster.

If the cluster is not compliant, the profile must be applied separately to each host within the cluster.

Appendixes

ESX Technical Support Commands

Most of the service console 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 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 vSphere 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 vSphere 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 vSphere Client or vCenter Server management.

The vSphere Client graphical user interface provides the preferred means of performing the configuration tasks described in this topic. You can use this topic to learn which vSphere Client commands to use in place of the service console commands. This topic provides a summary of the actions you take in vSphere Client but does not give complete instructions. For details on using commands and performing configuration tasks through vSphere Client, see the online help.

You can find additional information on a number of ESX commands by logging in 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, summarizes the purpose of each command, and provides a vSphere Client alternative. You can perform most of the vSphere Client actions listed in the table only after you have selected an ESX 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 Technical Support Commands

Service Console Command	Command Purpose and vSphere Client Procedure
<code>esxcfg-advcfg</code>	Configures advanced options for ESX. To configure advanced options in vSphere 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.
<code>esxcfg-auth</code>	Configures authentication. You can use this command to switch between the <code>pam_cracklib.so</code> and <code>pam_passwdc.so</code> plugins for password change rule enforcement. You also use this command to reset options for these two plugins. There is no means of configuring these functions in vSphere Client.
<code>esxcfg-boot</code>	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 vSphere Client.

Table A-1. ESX Technical Support Commands (Continued)

Service Console Command	Command Purpose and vSphere Client Procedure
<code>esxcfg-dumppart</code>	Configures a diagnostic partition or searches for existing diagnostic partitions. When you install ESX, 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 vSphere Client: ■ Determine whether there is a diagnostic partition — Click Storage>AddStorage 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 already has a diagnostic partition. ■ Configure a diagnostic partition — Click Storage>Add Storage>Diagnostic and step through the wizard.
<code>esxcfg-firewall</code>	Configures the service console firewall ports. To configure firewall ports for supported services and agents in vSphere Client, you select the Internet services that will be allowed to access the ESX host. Click Security Profile>Firewall>Properties and use the Firewall Properties dialog box to add services. You cannot configure unsupported services through the vSphere Client. For these services, use the <code>esxcfg-firewall</code>.
<code>esxcfg-info</code>	Prints information about the state of the service console, VMkernel, various subsystems in the virtual network, and storage resource hardware. vSphere 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.
<code>esxcfg-init</code>	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 host. There is no vSphere Client equivalent for this command.
<code>esxcfg-module</code>	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 vSphere Client equivalent for this command.
<code>esxcfg-mpath</code>	Configures multipath settings for your Fibre Channel or iSCSI disks. To configure multipath settings for your storage in vSphere 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.
<code>esxcfg-nas</code>	Manages NFS mounts. You use this command to create or unmount an NFS datastore. To view NFS datastores in vSphere Client, click Storage > Datastores and scroll through the datastores list. You can also perform the following activities from the Storage > Datastores view: ■ Display the attributes of an NFS datastore — Click the datastore and review the information under Details. ■ Create an NFS datastore — Click Add Storage. ■ Unmount an NFS datastore — Click Remove, or right-click the datastore to unmount and select Unmount.

Table A-1. ESX Technical Support Commands (Continued)

Service Console Command	Command Purpose and vSphere Client Procedure
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 vSphere Client, click Network Adapters. To change the speed and duplexing for a physical network adapter in the vSphere 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.
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 vSphere 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.
esxcfg-swiscsi	Configures your software iSCSI software adapter. To configure your software iSCSI system in vSphere 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.
esxcfg-upgrade	Upgrades from ESX Server 2.x to ESX. 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 vSphere Client: ■ Upgrade the host — You upgrade the binaries, converting from ESX Server 2.x to ESX. You cannot perform this step from vSphere Client. ■ 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-scsidevs	Prints a map of VMkernel storage devices to service console devices. There is no vSphere 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 vSphere 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.

Table A-1. ESX Technical Support Commands (Continued)

Service Console Command	Command Purpose and vSphere Client Procedure
<code>esxcfg-vswif</code>	Creates and updates service console network settings. This command is used if you cannot manage the ESX host through the vSphere Client because of network configuration issues. To set up connections for the service console in vSphere 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.
<code>esxcfg-vswitch</code>	Creates and updates virtual machine network settings. To set up connections for a virtual machine in vSphere 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.

Linux Commands Used with ESX

To support certain internal operations, ESX 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 functions unusable.

Always work through the vSphere Client when configuring ESX, unless otherwise instructed in vSphere documentation or by VMware Technical Support.

Using vmkfstools

You use the `vmkfstools` utility to create and manipulate virtual disks, file systems, logical volumes, and physical storage devices on the VMware ESX 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 vSphere Client.

This appendix includes the following topics:

- [“vmkfstools Command Syntax,”](#) on page 221
- [“vmkfstools Options,”](#) on page 222

vmkfstools Command Syntax

Generally, you do not 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 that 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 `vm1.<vm1_ID>:P` format, where `<vm1_ID>` is the device ID returned by the storage array and `P` is an integer that represents the partition number. The partition digit must be greater than zero (0) and should correspond to a valid VMFS partition of type `fb`.
- `<device>` specifies devices or logical volumes. This argument uses a path name in the ESX 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 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

The vmkfstools command has several options. Some of the options are suggested for advanced users only.

The long and single-letter forms of the options are equivalent. For example, the following commands are identical.

```
vmkfstools --createfs vmfs3 --blocksize 2m vml.<vml_ID>:1
vmkfstools -C vmfs3 -b 2m vml.<vml_ID>: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 is not 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 vSphere Client.

Creating a VMFS File System

Use the vmkfstools command to create a VMFS.

```
-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 vml.<vml_ID>:1. The partition becomes the file system's head partition.

VMFS-2 file systems are read-only on any ESX 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

This example illustrates creating a new VMFS-3 file system named `my_vmfs` on the `vm1.<vm1_ID>:1` partition. The file block size is 1MB.

```
vmkfstools -C vmfs3 -b 1m -S my_vmfs /vmfs/devices/disks/vml.<vm1_ID>:1
```

Extending an Existing VMFS-3 Volume

Use the `extendfs` command to add an extend to a VMFS volume.

```
-Z --extendfs <extension-device> <existing-VMFS-volume>
```

This command adds another extent to a previously created VMFS volume `<existing-VMFS-volume>`. You must specify the full path name, for example, `/vmfs/devices/disks/vml.<vm1_ID>:1`, not just the short name `vm1.<vm1_ID>: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 command, you lose all data that previously existed on the SCSI device you specified in `<extension-device>`.

Example for Extending a VMFS-3 Volume

This example extends the logical file system by allowing it to span to a new partition.

```
vmkfstools -Z /vmfs/devices/disks/vml.<vm1_ID_2>:1  
/vmfs/devices/disks/vml.<vm1_ID_1>:1
```

The extended file system spans two partitions: `vm1.<vm1_ID_1>:1` and `vm1.<vm1_ID_2>:1`. In this example, `vm1.<vm1_ID_1>:1` is the name of the head partition.

Listing Attributes of a VMFS Volume

The `queryfs` command lists the attributes of a VMFS volume.

```
-P --queryfs
    -h --human-readable
```

When you use this option on a 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 the 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, the capacity of the volume is listed 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 to 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 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 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 vSphere 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 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 host decides the allocation policy. The default allocation policy on most NFS servers is **thin**.

Creating a Virtual Disk

Use the **vmkfstools** command to create a virtual disk.

```
-c --createvirtualdisk <size>[kK|mM|gG]
    -a --adaptype [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.

Example for Creating a Virtual Disk

This example illustrates creating a 2 GB 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.

```
vmkfstools -c 2048m /vmfs/volumes/myVMFS/rh6.2.vmdk
```

Initializing a Virtual Disk

The `writezeros` command initializes 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, the initialization might take a long time.



CAUTION When you use this command, you lose all existing data on the virtual disk.

Inflating a Thin Virtual Disk

Use the `inflatedisk` command to inflate 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.

Deleting a Virtual Disk

The `deletevirtualdisk` option deletes files associated with the virtual disk listed at the specified path on the VMFS volume.

```
-U --deletevirtualdisk
```

Renaming a Virtual Disk

The `renamevirtualdisk` option renames a file associated with the virtual disk listed in the path specification portion of the command line. You must specify the original filename or file path `<oldName>` and the new filename or file path `<newName>`.

```
-E --renamevirtualdisk <oldName> <newName>
```

Cloning a Virtual or Raw Disk

You can create a copy of a virtual disk or raw disk.

```
-i --importfile <srcfile> -d --diskformat  
[rdm:<device>|rdmp:<device>|  
raw:<device>|thin|2gbsparse|monosparse|monoflat]
```

You can use the `-d` suboption for the `-i` option. This suboption specifies the disk format for the copy you create. A non-root user is not allowed to clone a virtual disk or a raw disk.

NOTE To clone the ESX Redo logs while preserving their hierarchy, use the `cp` command.

Example for Cloning a Virtual Disk

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.

```
vmkfstools -i /vmfs/volumes/templates/gold-master.vmdk /vmfs/volumes/myVMFS/myOS.vmdk
```

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
```

Migrate VMware Workstation and VMware GSX Server Virtual Machines

You cannot use a vSphere Client to migrate virtual machines created with VMware Workstation or VMware GSX Server into your ESX system. However, you can use the `vmkfstools -i` command to import the virtual disk into your ESX system and then attach this disk to a new virtual machine you create in ESX.

Prerequisites

You must import the virtual disk first because you cannot power on disks exported in `2gbsparse` format on an ESX host.

Procedure

- 1 Import a Workstation or GSX Server disk into your `/vmfs/volumes/myVMFS/` directory or any subdirectory.
- 2 In the vSphere 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 that you imported.

Extending a Virtual Disk

The `extendvirtualdisk` command extends the size of a disk allocated to a virtual machine after the virtual machine has been created.

```
-X --extendvirtualdisk <newSize>[kK|mM|gG]
```

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 that 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 by adding a `k` (kilobytes), `m` (megabytes), or `g` (gigabytes) suffix. The unit type is not case sensitive, so either `k` or `K` is interpreted as kilobytes. If you do not specify a unit type, the default is kilobytes.

The `newSize` parameter defines the entire new size, not just the increment you add to the disk.

For example, to extend a 4 GB virtual disk by 1 GB, enter `vmkfstools -X 5g <disk name>.disk`

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

The `migratevirtualdisk` option converts the specified virtual disk file from ESX Server 2 format to ESX format.

```
-M --migratevirtualdisk
```

Creating a Virtual Compatibility Mode Raw Device Mapping

The `createrdm` command 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.

```
-r --createrdm <device>
```

When specifying the `<device>` parameter, use the following format:

```
/vmfs/devices/disks/vml.<vml_ID>
```

NOTE All VMFS-3 file-locking mechanisms apply to RDMs.

Example for Creating a Virtual Compatibility Mode RDM

In this example, you create an RDM file named `my_rdm.vmdk` and map the `vml.<vml_ID>` raw disk to that file.

```
vmkfstools -r /vmfs/devices/disks/vml.<vml_ID> my_rdm.vmdk
```

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

The `createrdmpassthru` command maps a passthrough raw device to a file on a VMFS volume. This mapping lets a virtual machine bypass ESX 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.

```
-z --createrdmpassthru <device>
```

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, use the following format:

```
/vmfs/devices/disks/vml.<vml_ID>
```

Listing Attributes of an RDM

The `queryrdm` option lists the attributes of a raw disk mapping.

```
-q --queryrdm
```

This option prints the name of the RDM. The option also prints other identification information, like the disk ID, for the raw disk.

Displaying Virtual Disk Geometry

The `geometry` option gets information about the geometry of a virtual disk.

```
-g --geometry
```

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 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 tasks for physical storage devices. You can perform most of these tasks through the vSphere Client.

```
-L --lock [reserve|release|lunreset|targetreset|busreset]<device>
```

You can reserve a SCSI LUN for exclusive use by an ESX 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/vml.<vml_ID>:P
```


Index

Symbols

* next to path 111

Numerics

802.1Q and ISL tagging attacks 154

A

- accessing storage 78
- active adapters 26
- active uplinks 42, 45, 46
- active-active disk arrays 112
- active-passive disk arrays 112
- adapter, virtual 38
- adding
 - dvPort groups 32
 - NFS storage 97
- adding a VMkernel network adapter 22
- adding users to groups 167
- admin contact info 31
- Administrator role 164, 165
- aging, password restrictions 180
- antivirus software, installing 195
- applications
 - default 186, 187
 - disabling optional 186
 - optional 186, 187
 - setgid flag 185
 - setuid flag 185
- asterisk next to path 111
- attacks
 - 802.1Q and ISL tagging 154
 - double-encapsulated 154
 - MAC flooding 154
 - multicast brute-force 154
 - random frame 154
 - spanning tree 154
- authentication
 - groups 163
 - iSCSI storage 157
 - users 161, 163
 - vSphere Client to ESX 161
- authentication daemon 161
- average bandwidth 53, 54

B

- bandwidth
 - average 52, 53
 - peak 52, 53
- binding on host, dvPort groups 33
- Blade servers
 - and virtual networking 65
 - configuring a virtual machine port group 65
 - configuring a VMkernel port 66
- block devices 122
- blocked ports, dvPorts 55
- burst size 52–54

C

- CA-signed certificates 169
- CDP 27
- certificates
 - certificate file 168
 - checking 168
 - configuring host searches 171
 - default 168
 - disabling SSL for vSphere Web Access and SDK 171
 - generating new 169
 - key file 168
 - location 168
 - SSL 168
 - vCenter Server 168
 - vSphere Web Access 168
- certification, security 139
- changing host proxy services 172
- CHAP
 - disabling 92
 - for discovery targets 91
 - for iSCSI initiators 90
 - for static targets 91
 - mutual 89
 - one-way 89
- CHAP authentication 89, 157, 158
- CHAP authentication methods 89
- CIM and firewall ports 146
- cipher strength, connections 185
- Cisco Discovery Protocol 27, 31
- Cisco switches 27
- claim rules 110

- cloning raw disks **226**
- cloning virtual disks, example of **227**
- clusters, managing profiles **209**
- command reference for ESX **215**
- compatibility modes
 - physical **122**
 - virtual **122**
- config reset at disconnect, dvPort groups **33**
- configuring
 - dynamic discovery **88**
 - RDM **126**
 - SCSI storage **95**
 - static discovery **89**
- creating host profile **204, 205**
- creating host profiles **204**
- current multipathing state **111**
- cut and paste, disabling for guest operating systems **195**

D

- datastore copies, mounting **103**
- datastores
 - adding extents **101**
 - configuring on NFS volumes **97**
 - creating on SCSI disk **95**
 - displaying **81**
 - grouping **100**
 - increasing capacity **101**
 - managing **99**
 - managing duplicate **103**
 - mounting **103**
 - NFS **76**
 - paths **111**
 - refresh **94**
 - renaming **99**
 - review properties **82**
 - storage over-subscription **115**
 - unmounting **100**
 - viewing in vSphere Client **79**
 - VMFS **76**
- default certificates, replacing with CA-signed certificates **169**
- delegate user **95**
- deletevirtualdisk option **226**
- deployments for security
 - multiple customer open **191, 194**
 - multiple customer restricted **192**
- device disconnection, preventing **196**
- DHCP **25**
- diagnostic partition, configuring **98**
- direct access **163**
- disabling
 - cut and paste for virtual machines **195**

- iSCSI SAN authentication **158**
- logging for guest operating systems **198, 199**
- setgid applications **186**
- setuid applications **186**
- SSL for vSphere Web Access and SDK **171**
- variable information size **197**
- disabling paths **113**
- discovery
 - address **88**
 - dynamic **88**
 - static **89**
- disk arrays
 - active-active **112**
 - active-passive **112**
- disk formats
 - NFS **95**
 - supported **225**
 - thick provisioned **114**
 - thin provisioned **114**
- disks, format **115**
- DMZ **136**
- DNS **55**
- double-encapsulated attacks **154**
- dvPort group, load balancing **45**
- dvPort groups
 - binding on host **33**
 - config reset at disconnect **33**
 - description **32**
 - failback **45**
 - failover order **45**
 - live port moving **33**
 - name **32**
 - network failover detection **45**
 - notify switches **45**
 - number of ports **32**
 - override settings **33**
 - port blocking **54**
 - port group type **32**
 - port name format **33**
 - teaming and failover policies **45**
 - traffic shaping policies **53**
 - virtual machines **40**
- dvPort Groups, adding **32**
- dvPorts
 - blocked ports **55**
 - failback **46**
 - failover order **46**
 - load balancing **46**
 - network failover detection **46**
 - notify switches **46**
 - port policies **55**
 - properties **33**
 - teaming and failover policies **46**

- traffic shaping policies **54**
- VLAN policies **48**
- DVS
 - adding a VMkernel network adapter **36**
 - admin contact info **31**
 - Cisco Discovery Protocol **31**
 - IP address **31**
 - maximum MTU **31**
 - maximum number of ports **31**
- dvUplink **30**
- dynamic discovery, configuring **88**
- dynamic discovery addresses **88**
- E**
 - early binding port groups **32**
 - educational support **9**
 - encryption
 - certificates **168**
 - enabling and disabling SSL **168**
 - for user name, passwords, packets **168**
 - enhanced vmxnet **57, 59**
 - ESX
 - command reference **215**
 - using Linux commands **219**
 - esxcfg commands **215**
 - exporting
 - host groups **165**
 - host users **165**
 - extendfs option **223**
 - extendvirtualdisk command **227**
 - extents
 - adding to datastore **101**
 - growing **101**
- F**
 - failback **42, 45, 46**
 - failover **42, 105**
 - failover order **42, 45, 46**
 - failover paths, status **111**
 - failover policies
 - dvPort groups **45**
 - dvPorts **46**
 - port groups **44**
 - vSwitch **42**
 - Fibre Channel **72**
 - Fibre Channel SANs, WWNs **74**
 - Fibre Channel storage, overview **84**
 - file systems, upgrading **102**
 - firewall ports
 - automating service behavior **148**
 - backup agents **176**
 - closing **179**
 - configuring with vCenter Server **142**
 - configuring without vCenter Server **143**
 - connecting to vCenter Server **144**
 - connecting virtual machine console **145**
 - encryption **168**
 - host to host **146**
 - management **146**
 - opening in service console **178**
 - opening with vSphere Client **146**
 - overview **141**
 - SDK and virtual machine console **145**
 - security level **176, 177**
 - service console **176–179**
 - supported services **146**
 - vSphere Client and vCenter Server **142**
 - vSphere Client and virtual machine console **145**
 - vSphere Client direct connection **143**
 - vSphere Web Access and the virtual machine console **145**
 - vSphere Web Access and vCenter Server **142**
 - vSphere Web Access direct connection **143**
 - firewalls
 - access for management agents **147**
 - access for services **147**
 - configuring **148**
 - Fixed path policy **106, 112**
 - forged transmissions **155, 156**
 - forged transmits **50, 51**
 - FTP and firewall ports **146**
- G**
 - gateways, setting the default **24**
 - generating certificates **169**
 - geometry option **228**
 - groups
 - about **165**
 - adding to hosts **167**
 - adding users **167**
 - authentication **163**
 - exporting a group list **165**
 - modifying on hosts **167**
 - permissions and roles **162**
 - removing from hosts **167**
 - viewing group lists **165**
 - guest operating systems
 - disabling cut and paste **195**
 - disabling logging **198, 199**
 - limiting variable information size **197**
 - logging levels **198**
 - security recommendations **195**

H

- hardware devices, removing **196**
- hardware iSCSI, and failover **108**
- hardware iSCSI initiator, changing iSCSI name **86**
- hardware iSCSI initiators
 - configuring **85**
 - installing **85**
 - setting up discovery addresses **88**
 - setting up naming parameters **86**
 - viewing **85**
- host certificate searches **171**
- host networking, viewing **16**
- host profile, attaching entities **207**
- host profiles
 - accessing **204**
 - applying profiles **208**
 - attaching entities from host **207**
 - attaching entities from Host Profile view **207**
 - changing the reference host **209**
 - checking compliance **210**
 - create new profile **204**
 - create new profile from host **205**
 - create new profile from host profile view **204**
 - editing **206**
 - editing policies **206**
 - enabling policy compliance check **207**
 - exporting **205**
 - importing **205**
 - managing profiles **207**
 - usage model **203**
- host-to-host firewall ports **146**
- hosts
 - adding groups **167**
 - adding to a vNetwork Distributed Switch **30**
 - adding users **166**
 - deployments and security **191**
 - memory **197**
 - thumbprints **168**

I

- IDE **72**
- inbound traffic shaping **53, 54**
- inflatedisk command **226**
- initializing virtual disks **226**
- Intel Nehalem platforms **60**
- Internet Protocol **41**
- IP address **31**
- IP addresses **74**
- IP storage port groups, creating **22, 36**
- IPv4 **41**
- IPv6 **41**

iSCSI

- authentication **157**
- networking **64**
- protecting transmitted data **158**
- QLogic iSCSI adapters **157**
- securing ports **158**
- security **157**
- software client and firewall ports **146**
- iSCSI aliases **74**
- iSCSI HBA, alias **86**
- iSCSI initiators
 - advanced parameters **93**
 - configuring advanced parameters **93**
 - configuring CHAP **90**
 - hardware **85**
 - setting up CHAP parameters **89**
- iSCSI names **74**
- iSCSI networking, creating a VMkernel port **64, 87**
- iSCSI SAN authentication, disabling **158**
- iSCSI storage
 - hardware-initiated **84**
 - initiators **84**
 - software-initiated **84**
- isolation
 - virtual machines **132**
 - virtual networking layer **134**
 - virtual switches **134**
 - VLANs **134**

J

- jumbo frames
 - enabling **58, 59**
 - virtual machines **57, 59**

L

- late binding port groups **32**
- Layer 2 security **48**
- Linux commands, using with ESX **219**
- live port moving, dvPort groups **33**
- load balancing **42, 44–46**
- local SCSI storage, overview **83**
- log files
 - limiting number **199**
 - limiting size **199**
- logging, disabling for guest operating systems **198, 199**
- logging levels, guest operating systems **198**
- logical file systems, example of extending **223**
- LUNs
 - creating, and rescan **94**
 - making changes and rescan **94**
 - managing paths **127**
 - managing SCSI reservations **229**

- masking changes and rescan **94**
- multipathing policy **112**
- setting multipathing policy **112**

M

- MAC address
 - configuring **55**
 - generating **56**
- MAC address changes **155, 156**
- MAC addresses, adding **56**
- MAC flooding **154**
- management access
 - firewalls **147**
 - TCP and UDP ports **149**
- managing paths, LUNs **127**
- mapped raw LUNs **127**
- maximum MTU **31**
- maximum number of ports **31**
- metadata, RDMS **122**
- migratevirtualdisk option **227**
- modifying groups on hosts **167**
- Most Recently Used path policy **106, 112**
- mounting VMFS datastores **103**
- MPPs, See multipathing plugins
- MRU path policy **112**
- MTU **58**
- multicast brute-force attacks **154**
- multipathing
 - active paths **111**
 - broken paths **111**
 - disabled paths **111**
 - standby paths **111**
 - viewing the current state of **111**
- multipathing plugins, path claiming **110**
- multipathing policy **112**
- multipathing state **111**
- mutual CHAP **89**

N

- NAS, mounting **64**
- NAT **41**
- Native Multipathing Plugin **105, 106**
- Nessus **188**
- NetQueue
 - disabling **60**
 - enabling **60**
- network adapter, service console **23**
- network adapters, viewing **17, 31**
- network address translation **41**
- network failover detection **42, 44–46**
- networking, security policies **50, 51**
- networking best practices **63**
- networking performance, NetQueue **60**

- networking policies **42**
- networking stack **21**
- networks, security **150**
- NFS, firewall ports **146**
- NFS datastores
 - as repositories **96**
 - unmounting **100**
- NFS storage
 - adding **97**
 - overview **95**
- NIC, adding **35**
- NIC teaming, definition **15**
- NIS and firewall ports **146**
- NMP, path claiming **110**
- no access role **165**
- No Access role **164**
- notify switches **42, 45, 46**
- NTP **148**

O

- one-way CHAP **89**
- outbound traffic shaping **53, 54**
- override settings, dvPort groups **33**

P

- pam_cracklib.so plug-in **181**
- pam_passwdqc.so plug-in **184**
- partition mappings **122**
- passive disk arrays **112**
- passthrough device, add to a virtual machine **61**
- passthrough devices
 - configuring **61**
 - mapping **228**
- passwords
 - aging **180**
 - aging restrictions **180**
 - complexity **181, 183**
 - criteria **181**
 - host **179–184**
 - length **181**
 - pam_cracklib.so plug-in **181**
 - pam_passwdqc.so plug-in **184**
 - plug-ins **181**
 - restrictions **179–181**
 - reuse rules **182**
 - service console **179**
- path claiming **110**
- path failover, host-based **108**
- path failure **107**
- path failure rescan **94**
- path management **105**
- path policies
 - changing defaults **112**

- Fixed **106, 112**
- Most Recently Used **106, 112**
- MRU **112**
- Round Robin **106, 112**
- Path Selection Plugins **106**
- paths
 - disabling **113**
 - preferred **111**
- peak bandwidth **52–54**
- permissions
 - and privileges **163**
 - overview **163**
 - root user **163**
 - user **164**
 - vCenter Server administrator **163**
 - vpxuser **163**
- physical adapters, removing **36**
- physical switches, troubleshooting **68**
- plug-ins
 - pam_cracklib.so **181**
 - pam_passwdqc.so **184**
- Pluggable Storage Architecture **105**
- port binding **87, 108**
- port blocking, dvPort groups **54**
- port configuration **25**
- port group
 - definition **15**
 - using **20**
- port groups
 - failback **44**
 - failover order **44**
 - Layer 2 Security **49**
 - load balancing **44**
 - network failover detection **44**
 - renaming **68**
 - teaming and failover policies **44**
 - traffic shaping **53**
 - troubleshooting **68**
- port name format, dvPort groups **33**
- ports, service console **24**
- preferred path **111**
- private VLAN
 - create **34**
 - primary **34**
 - removing **34, 35**
 - secondary **35**
- private VLANs **34**
- privileges and permissions **163**
- profiles, managing from a cluster **209**
- promiscuous mode **50, 51, 155, 156**
- properties, dvPorts **33**

- proxy services
 - changing **172**
 - encryption **168**
- PSA, See Pluggable Storage Architecture
- PSPs, See Path Selection Plugins

Q

- queryfs command **224**

R

- RAID devices **122**
- random frame attacks **154**
- raw device mapping, see RDM **119**
- raw disk mappings, listing attributes of **228**
- raw disks, cloning **226**
- RDM
 - advantages **120**
 - and virtual disk files **125**
 - creating **126**
 - dynamic name resolution **123**
 - overview **119**
 - physical compatibility mode **122**
 - virtual compatibility mode **122**
 - with clustering **125**
- RDM files, creating **228**
- RDMs
 - and snapshots **122**
 - and VMFS formats **122**
 - example of creating **228**
 - listing attributes of **228**
- Read Only role **164, 165**
- reference hosts, changing **209**
- refresh **94**
- removing users from groups **167**
- renamevirtualdisk option **226**
- replacing, default certificates **169**
- repositories, NFS datastores **96**
- rescan
 - LUN creation **94**
 - path masking **94**
 - when path is down **94**
- resource limits and guarantees, security **132**
- roles
 - Administrator **164**
 - and permissions **164**
 - default **164**
 - No Access **164**
 - Read Only **164**
 - security **164**
- root login
 - permissions **163, 164**
 - SSH **187**
- Round Robin path policy **106, 112**
- routing **55**

S

- SAS **72**
- SATA **72**
- SCSI, vmkfstools **221**
- SCSI reservations, managing **229**
- SDK, firewall ports and virtual machine console **145**
- security
 - architecture **131**
 - certification **139**
 - cipher strength **185**
 - DMZ in single host **134, 136**
 - features **131**
 - iSCSI storage **157**
 - overview **131**
 - PAM authentication **161**
 - patches **188**
 - permissions **163**
 - recommendations for virtual machines **195**
 - resource guarantees and limits **132**
 - scanning software **188**
 - service console **138, 175**
 - setuid and setgid flags **185**
 - virtual machines **132**
 - virtual machines with VLANs **150**
 - virtual networking layer **134**
 - virtual switch ports **155**
 - virtualization layer **132**
 - VLAN hopping **152**
 - VMkernel **132**
 - VMware policy **139**
 - vmware-authd **161**
 - vmware-hostd **161**
- security policies, dvPorts **50, 51**
- service console
 - closing firewall ports **179**
 - direct connections **176**
 - firewall ports **177**
 - firewall security **176**
 - isolating **153**
 - logging in **176**
 - network policies **25**
 - opening firewall ports **178**
 - password plug-in **184**
 - password restrictions **179**
 - recommendations for securing **175**
 - remote connections **176**
 - securing with VLANs and virtual switches **152**
 - security **138**
 - setgid applications **185**
 - setuid applications **185**
 - SSH connections **187**
 - troubleshooting **67**
 - VLAN **25**
- service console configuration, editing **39**
- service console network adapters **37**
- service console networking
 - configuration **23**
 - troubleshooting **67**
- services
 - automating **148**
 - starting **148**
- setgid
 - applications **185**
 - default applications **187**
 - disabling applications **186**
- setinfo **197**
- setuid
 - applications **185**
 - default applications **186**
 - disabling applications **186**
- shell access, granting **166**
- single point of failure **83**
- SMB and firewall ports **146**
- SNMP and firewall ports **146**
- software iSCSI
 - and failover **108**
 - diagnostic partition **97**
 - networking **87**
- software iSCSI initiators
 - configuring **86**
 - enabling **87**
 - setting up discovery addresses **88**
- spanning tree attacks **154**
- SPOF **83**
- SSH
 - configuring **188**
 - firewall ports **146**
 - security settings **187**
 - service console **187**
- SSL
 - enabling and disabling **168**
 - encryption and certificates **168**
 - timeouts **170**
- standby adapters **26**
- standby uplinks **42, 45, 46**
- static discovery, configuring **89**
- static discovery addresses **88**
- storage
 - access for virtual machines **78**
 - adapters **73**
 - Fibre Channel **84**
 - iSCSI **84**
 - local **72**
 - local SCSI **83**

- networked **72**
- NFS **95**
- not-shared **114**
- overview **71**
- provisioned **114**
- provisioning **113**
- SAN **84**
- securing with VLANs and virtual switches **152**
- types **72**
- used by virtual machines **114**
- viewing in vSphere Client **79**
- storage adapters
 - copying names **80**
 - Fibre Channel **84**
 - viewing **79**
 - viewing in vSphere Client **79**
- SATPs **106**
- Storage Array Type Plugins **106**
- storage device identifiers, copying **81**
- storage devices
 - displaying for a host **80**
 - displaying for an adapter **81**
 - identifiers **75**
 - names **75**
 - paths **111**
 - runtime names **75**
 - viewing **80**
- storage filters, turning off **116**
- storage space **113**
- storage types, comparing **78**
- switch, vNetwork **38**

T

- targets **73**
- TCP ports **149**
- TCP/IP stack, setting default gateway **24**
- teaming policies
 - dvPort groups **45**
 - dvPorts **46**
 - vSwitch **42**
- technical support **9**
- thin disks, creating **114**
- thin virtual disks, inflating **226**
- third-party software support policy **139**
- third-party switch **29**
- thumbprints, hosts **168**
- timeouts, SSL **170**
- Tomcat Web service **138**
- traffic shaping
 - port groups **53**
 - vSwitch **52**

- traffic shaping policies
 - dvPort groups **53**
 - dvPorts **54**
- TSO, checking whether enabled **58**

U

- UDP ports **149**
- updated information **7**
- uplink adapters
 - adding **26**
 - duplex **26**
 - speed **26**
- uplink assignments **31**
- uplinks, removing **36**
- USB **72**
- user permissions, vpxuser **164**
- user roles
 - Administrator **165**
 - no access **165**
 - Read Only **165**
- users
 - about **165**
 - adding to groups **167**
 - adding to hosts **166**
 - authentication **163**
 - direct access **163**
 - exporting a user list **165**
 - from Windows domain **163**
 - modifying on hosts **166**
 - permissions and roles **162**
 - removing from groups **167**
 - removing from hosts **167**
 - security **163**
 - vCenter Server **163**
 - viewing user list **165**

V

- variable information size for guest operating systems
 - disabling **197**
 - limiting **197**
- vCenter Server
 - connecting through firewall **144**
 - firewall ports **142**
 - permissions **163**
- vCenter Server users **163**
- verbose option **222**
- virtual adapter, VMkernel **38**
- virtual adapters, migrating **37**
- virtual disks
 - cloning **226**
 - creating **225**
 - deleting **226**
 - example of cloning **227**

- example of creating **226**
- extending **227**
- formats **114**
- getting geometry information **228**
- inflating **226**
- initializing **226**
- migrating **227**
- renaming **226**
- supported disk formats **225**
- virtual machine networking **16, 20**
- virtual machines
 - creating with RDMs **126**
 - disabling cut and paste **195**
 - disabling logging **198, 199**
 - isolation **134, 136**
 - limiting variable information size **197**
 - migrating **227**
 - migrating to or from a vNetwork Distributed Switch **40**
 - networking **40**
 - preventing device disconnection **196**
 - resource reservations and limits **132**
 - security **132**
 - security recommendations **195**
- virtual network, security **150**
- virtual network adapters, removing **39**
- virtual networking layer and security **134**
- virtual switch ports, security **155**
- virtual switch security **152**
- virtual switches
 - 802.1Q and ISL tagging attacks **154**
 - and iSCSI **158**
 - double-encapsulated attacks **154**
 - forged transmissions **155**
 - MAC address changes **155**
 - MAC flooding **154**
 - multicast brute-force attacks **154**
 - promiscuous mode **155**
 - random frame attacks **154**
 - scenarios for deployment **191**
 - security **154**
 - spanning tree attacks **154**
- virtualization layer, security **132**
- VLAN
 - definition **15**
 - private **34**
- VLAN ID **32**
- VLAN policies
 - dvPort group **48**
 - dvPorts **48**
- VLAN security **152**
- VLAN trunking **48**
- VLAN Trunking **32, 48**
- VLAN Type **48**
- VLANs
 - and iSCSI **158**
 - Layer 2 security **152**
 - scenarios for deployment **191**
 - security **150, 153**
 - VLAN hopping **152**
- VLANS
 - configuring for security **153**
 - service console **153**
- VMDirectPatch **60**
- VMFS
 - creating **222**
 - example of creating **223**
 - sharing **191**
 - vmkfstools **221**
 - volume resignaturing **103**
- VMFS datastores
 - adding extents **101**
 - changing properties **101**
 - changing signatures **104**
 - configuring **95**
 - creating **76**
 - deleting **100**
 - increasing capacity **101**
 - resignaturing copies **104**
 - sharing **77**
 - unmounting **100**
- VMFS volume, extending **223**
- VMFS volume resignaturing **103**
- VMFS volumes, listing attributes of **224**
- VMFS-2 volumes
 - migrating to VMFS-3 **227**
 - upgrading to VMFS-3 **224**
- VMFS-3 volume, example of extending **223**
- VMFS-3 volumes, upgrading from VMFS-2 **224**
- VMkernel, TCP/IP stack **21**
- VMkernel
 - configuring **21**
 - definition **15**
 - security **132**
- VMkernel adapter **38**
- VMkernel network adapters, adding **22, 36**
- VMkernel networking **16**
- vmkfstools
 - file system options **222**
 - overview **221**
 - syntax **221**
 - virtual disk options **225**
- vmkfstools command
 - options **222**
 - verbose option **222**
- vmkfstools options **222**

- VMotion
 - definition **15**
 - networking configuration **21**
 - securing with VLANs and virtual switches **152**
- VMotion interfaces, creating **22, 36**
- VMware NMP
 - I/O flow **107**
 - See also* Native Multipathing Plugin
- vmware-hostd **161**
- vNetwork Distributed Switch
 - adding a host to **30**
 - adding a NIC to **35**
 - new **30**
 - third-party **29**
 - VMkernel adapter **38**
- vNetwork Distributed Switches
 - adding a VMkernel network adapter **36**
 - adding hosts to **30**
 - admin contact info **31**
 - Cisco Discovery Protocol **31**
 - connecting virtual machines **39**
 - creating service console network adapters **37**
 - general settings, editing **31**
 - IP address **31**
 - maximum MTU **31**
 - maximum number of ports **31**
 - migrating virtual adapters **37**
 - migrating virtual machines to or from **40**
 - miscellaneous policies **55**
 - physical adapters **35**
 - service console configuration, editing **39**
- vNetwork Standard Switch
 - Layer 2 security **49**
 - port configuration **25**
 - traffic shaping **52**
 - viewing **16**
- vNetwork Standard Switches, migrating virtual adapters **37**
- volume resignaturing **103, 104**
- vpxuser **164**
- vSphere Client
 - firewall ports connecting to virtual machine console **145**
 - firewall ports for direct connection **143**
 - firewall ports with vCenter Server **142**
 - migrating virtual machines **227**
- vSphere Web Access
 - and host services **168**
 - disabling SSL **171**
 - firewall ports connecting to virtual machine console **145**
 - firewall ports for direct connection **143**
 - firewall ports with vCenter Server **142**
- vSwitch
 - definition **15**
 - failback **42**
 - failover order **42**
 - Layer 2 security **49**
 - load balancing **42**
 - network failover detection **42**
 - notify switches **42**
 - port configuration **25**
 - teaming and failover policies **42**
 - traffic shaping **52**
 - using **19**
 - viewing **16**

W

- writezeros command **226**
- WWNs **74**