

# vSphere Networking

ESXi 5.0  
vCenter Server 5.0

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-000599-01

**vmware**<sup>®</sup>

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](mailto: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](http://www.vmware.com)

# Contents

	About vSphere Networking	5
<b>1</b>	<b>Updated Information</b>	<b>7</b>
<b>2</b>	<b>Introduction to Networking</b>	<b>9</b>
	Networking Concepts Overview	9
	Network Services	10
	View Networking Information in the vSphere Client	10
	View Network Adapter Information in the vSphere Client	11
<b>3</b>	<b>Setting Up Networking with vSphere Standard Switches</b>	<b>13</b>
	vSphere Standard Switches	13
	Standard Port Groups	14
	Port Group Configuration for Virtual Machines	14
	VMkernel Networking Configuration	15
	vSphere Standard Switch Properties	18
<b>4</b>	<b>Setting Up Networking with vSphere Distributed Switches</b>	<b>21</b>
	vSphere Distributed Switch Architecture	22
	Configuring a vSphere Distributed Switch	22
	Distributed Port Groups	27
	Working with Distributed Ports	28
	Private VLANs	29
	Configuring vSphere Distributed Switch Network Adapters	31
	Configuring Virtual Machine Networking on a vSphere Distributed Switch	35
<b>5</b>	<b>Managing Network Resources</b>	<b>37</b>
	vSphere Network I/O Control	37
	TCP Segmentation Offload and Jumbo Frames	40
	NetQueue and Networking Performance	42
	DirectPath I/O	43
<b>6</b>	<b>Networking Policies</b>	<b>45</b>
	Load Balancing and Failover Policy	45
	VLAN Policy	52
	Security Policy	52
	Traffic Shaping Policy	56
	Resource Allocation Policy	59
	Monitoring Policy	60
	Port Blocking Policies	61
	Manage Policies for Multiple Port Groups on a vSphere Distributed Switch	62

<b>7</b>	<b>Advanced Networking</b>	<b>67</b>
	Enable Internet Protocol Version 6 Support	67
	VLAN Configuration	68
	Working With Port Mirroring	68
	Configure NetFlow Settings	72
	Switch Discovery Protocol	72
	Change the DNS and Routing Configuration	74
	MAC Addresses	74
	Mounting NFS Volumes	76
<b>8</b>	<b>Networking Best Practices</b>	<b>77</b>
	Index	79

# About vSphere Networking

---

*vSphere Networking* provides information about configuring networking for VMware vSphere<sup>®</sup>, including how to create vSphere distributed switches and vSphere standard switches.

*vSphere Networking* also provides information on monitoring networks, managing network resources, and networking best practices.

## Intended Audience

The information presented is written for experienced Windows or Linux system administrators who are familiar with network configuration and virtual machine technology.



## Updated Information

---

This *vSphere Networking* documentation is updated with each release of the product or when necessary.

This table provides the update history of *vSphere Networking*.

Revision	Description
EN-000599-01	<ul style="list-style-type: none"><li>■ Added clarification to the section <a href="#">“DirectPath I/O,”</a> on page 43 to describe vMotion requirements and functionality using Cisco-specific switches.</li><li>■ Added clarification to the section <a href="#">“Enable DirectPath I/O with vMotion on a Virtual Machine,”</a> on page 44 to describe Cisco-specific switch functionality with vMotion.</li><li>■ Added the section <a href="#">“Removing NICs from Active Virtual Machines,”</a> on page 32 to describe the behavior of the vSphere Client when NICs are removed from active virtual machines.</li></ul>
EN-000599-00	Initial release.



# Introduction to Networking

---

The basic concepts of ESXi networking and how to set up and configure a network in a vSphere environment are discussed.

This chapter includes the following topics:

- [“Networking Concepts Overview,”](#) on page 9
- [“Network Services,”](#) on page 10
- [“View Networking Information in the vSphere Client,”](#) on page 10
- [“View Network Adapter Information in the vSphere Client,”](#) on page 11

## Networking Concepts Overview

A few concepts are essential for a thorough understanding of virtual networking. If you are new to ESXi, 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 ESXi 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 vSphere standard switch 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 vSphere standard switch 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 vSphere standard switch works much like a physical switch, it does not have some of the advanced functionality of a physical switch.

A vSphere distributed switch acts as a single switch across all associated hosts on a datacenter. This allows virtual machines to maintain consistent network configuration as they migrate across multiple hosts.

A distributed port is a port on a vSphere distributed switch that connects to a host’s 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 standard switches through port groups. Port groups define how a connection is made through the switch to the network. Typically, a single standard switch is associated with one or more port groups.

A distributed port group is a port group associated with a vSphere distributed switch and specifies port configuration options for each member port. Distributed port groups define how a connection is made through the vSphere distributed switch to the network.

NIC teaming occurs when multiple uplink adapters are associated with a single switch 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, vMotion, and Fault Tolerance Logging. Virtual machines run their own systems' TCP/IP stacks and connect to the VMkernel at the Ethernet level through standard and distributed 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 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 ESXi host to another without shutting down the virtual machine. The optional vMotion feature requires its own license key.

## Network Services

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

You can to enable two types of network services in ESXi:

- Connecting virtual machines to the physical network and to each other.
- Connecting VMkernel services (such as NFS, iSCSI, or vMotion) to the physical network.

## View Networking Information in the vSphere Client

The vSphere Client shows 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 (Optional) Choose the type of networking to view.

Option	Description
<b>vSphere Standard Switch</b>	Displays vSphere standard switch networking on the host.
<b>vSphere Distributed Switch</b>	Displays vSphere distributed switch networking on the host.

The **vSphere Distributed Switch** option appears only on hosts that are connected to one or more vSphere distributed switches.

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 **Hosts and Clusters** inventory view.
- 2 Select the host in the inventory pane.
- 3 Click the **Configuration** tab, and click **Network Adapters**.

The network adapters panel shows the following information.

**Table 2-1.** Network Adapter Parameters

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.
Switch	vSphere standard switch or vSphere distributed switch that the network adapter is associated with.
Observed IP ranges	IP addresses that the network adapter is likely to have access to.
Wake on LAN supported	Network adapter ability to support Wake on the LAN.



# Setting Up Networking with vSphere Standard Switches

---

# 3

vSphere standard switches handle network traffic at the host level in a vSphere environment.

Use the vSphere Client to add networking based on the categories that reflect the types of network services.

- Virtual machines
- VMkernel

This chapter includes the following topics:

- [“vSphere Standard Switches,”](#) on page 13
- [“Standard Port Groups,”](#) on page 14
- [“Port Group Configuration for Virtual Machines,”](#) on page 14
- [“VMkernel Networking Configuration,”](#) on page 15
- [“vSphere Standard Switch Properties,”](#) on page 18

## vSphere Standard Switches

You can create abstracted network devices called vSphere standard switches. A standard switch can route traffic internally between virtual machines and link to external networks.

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

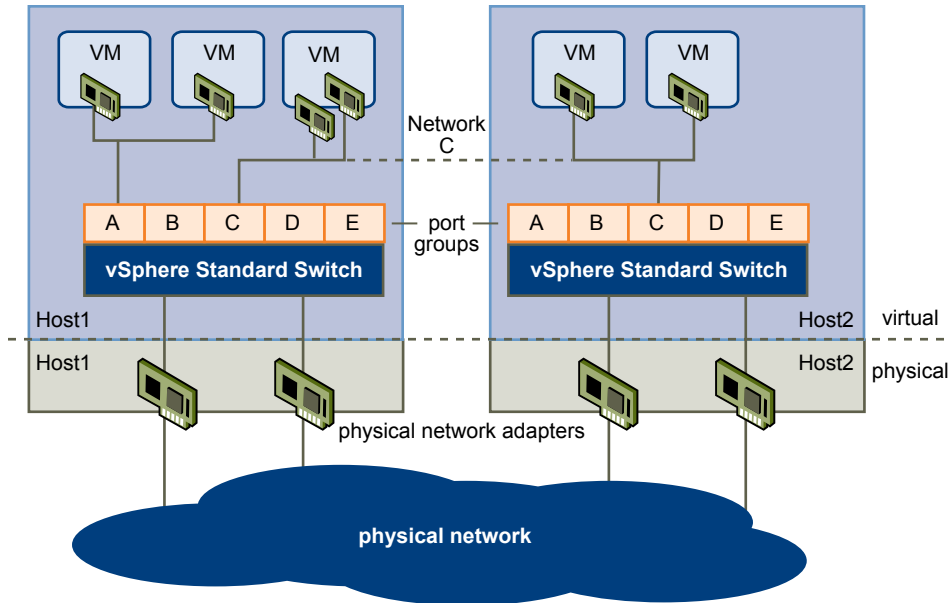
A vSphere standard switch models a physical Ethernet switch. The default number of logical ports for a standard switch is 120. You can connect one network adapter of a virtual machine to each port. Each uplink adapter associated with a standard switch uses one port. Each logical port on the standard switch is a member of a single port group. Each standard switch can also have one or more port groups assigned to it. For information about maximum allowed ports and port groups, see the *Configuration Maximums* documentation.

When two or more virtual machines are connected to the same standard switch, network traffic between them is routed locally. If an uplink adapter is attached to the standard switch, each virtual machine can access the external network that the adapter is connected to.

## Standard Port Groups

Port groups aggregate multiple ports under a common configuration and provide a stable anchor point for virtual machines connecting to labeled networks.

**Figure 3-1.** vSphere Standard Switch Network



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 vSphere standard switch 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.

ESXi 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 vSphere standard switch 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 standard switch, 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 Select the host in the inventory pane.

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

4 Select the vSphere Standard Switch view.

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

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

6 Accept the default connection type, **Virtual Machines**, and click **Next**.

7 Select **Create a vSphere standard switch** or one of the listed existing standard switches and the associated physical adapters to use for this port group.

You can create a new standard switch with or without Ethernet adapters.

If you create a standard switch without physical network adapters, all traffic on that switch is confined to that switch. No other hosts on the physical network or virtual machines on other standard switches can send or receive traffic over this standard switch. You might create a standard switch 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.

8 Click **Next**.

9 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.

10 (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.

11 Click **Next**.

12 After you determine that the switch is configured correctly, click **Finish**.

## VMkernel Networking Configuration

A VMkernel networking interface provides network connectivity for the host as well as handling VMware vMotion, IP storage, and Fault Tolerance.

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 ESXi. Because these storage types are network based, they can use the same VMkernel interface and port group.

## 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.
- Fault Tolerance logging.
- Port-binding for vMotion interfaces.
- Provides networking information to dependent hardware iSCSI adapters.

If you have two or more physical NICs for iSCSI, you can create multiple paths for the software iSCSI by configuring iSCSI Multipathing. For more information about iSCSI Multipathing, see the *vSphere Storage* documentation.

---

**NOTE** ESXi supports only NFS version 3 over TCP/IP.

---

## Set Up VMkernel Networking on a vSphere Standard 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 **Hosts and Clusters** inventory view.
- 2 Select the host in the inventory pane.
- 3 On the host **Configuration** tab, click **Networking**.
- 4 In the vSphere Standard Switch view, click **Add Networking**.
- 5 Select **VMkernel** and click **Next**.
- 6 Select the vSphere standard switch to use, or select **Create a vSphere standard switch** to create a new vSphere standard switch.
- 7 Select the check boxes for the network adapters for your vSphere standard switch to use.  
 Select adapters for each vSphere standard switch 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 vSphere standard switch, all the network adapters in the system are being used by existing vSphere standard switches or vSphere distributed switches. You can either create a vSphere standard switch without a network adapter, or select a network adapter that an existing vSphere standard switch uses.
- 8 Click **Next**.

- 9 Select or enter a network label and a VLAN ID.

Option	Description
<b>Network Label</b>	A name that identifies the port group that you are creating. This is the label that you specify when you configure VMkernel services such as vMotion and IP storage and you configure a virtual adapter to be attached to this port group.
<b>VLAN ID</b>	Identifies the VLAN that the port group's network traffic will use.

- 10 (Optional) Select **Use this port group for vMotion** to enable this port group to advertise itself to another host as the network connection through which vMotion traffic should be sent.
- 11 (Optional) Select **Use this port group for fault tolerance logging**.
- 12 (Optional) Select **Use this port group for management traffic**.
- 13 If IPv6 is enabled on the host, select **IP (Default)**, **IPv6**, or **IP and IPv6 networking**.  
This option does not appear on hosts that do not have IPv6 enabled. IPv6 configuration cannot be used with dependent hardware iSCSI adapters.
- 14 Click **Next**.
- 15 Select how to obtain IP settings.

Option	Description
<b>Obtain IP settings automatically</b>	Use DHCP to obtain IP settings.
<b>Use the following IP settings</b>	Specify IP settings manually. <ol style="list-style-type: none"> <li>Enter the IP address and subnet mask for the VMkernel interface.</li> <li>Click <b>Edit</b> to set the VMkernel Default Gateway for VMkernel services, such as vMotion, NAS, and iSCSI.  On the <b>DNS Configuration</b> 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.</li> <li>Click <b>OK</b> and click <b>Next</b>.</li> </ol>

- 16 If you are using IPv6 for the VMkernel interface, select an option for obtaining IPv6 addresses.

Option	Description
<b>Obtain IPv6 addresses automatically through DHCP</b>	Use DHCP to obtain IPv6 addresses.
<b>Obtain IPv6 addresses automatically through router advertisement</b>	Use router advertisement to obtain IPv6 addresses.
<b>Static IPv6 addresses</b>	<ol style="list-style-type: none"> <li>Click <b>Add</b> to add a new IPv6 address.</li> <li>Enter the IPv6 address and subnet prefix length, and click <b>OK</b>.</li> <li>To change the VMkernel default gateway, click <b>Edit</b>.</li> </ol>

- 17 Click **Next**.
- 18 Review the information, click **Back** to change any entries, and click **Finish**.

## View VMkernel Routing Information on a vSphere Standard Switch

You can view IP and IPv6 routing information, such as network, prefix, and gateway, for a VMkernel network interface on a vSphere standard switch.

### Procedure

- 1 Log in to the vSphere Client and select the **Hosts and Clusters** inventory view.

- 2 On the host **Configuration** tab, click **Networking**.
- 3 Click **Properties** for the standard switch associated with the VMkernel interface to view.
- 4 On the **Ports** tab, select the VMkernel network adapter to view, and click **View Routing Table** under IP Settings or IPv6 Settings.

A routing table that includes network, prefix, and gateway information for the selected VMkernel network adapter appears.

## vSphere Standard Switch Properties

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

### Change the Number of Ports for a vSphere Standard Switch

A vSphere standard switch 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 standard switch that you want to edit.
- 4 Click the **Ports** tab.
- 5 Select the standard switch 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 standard switch 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 vSphere standard switch 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 standard switch 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
<b>Active Adapters</b>	Adapters that the standard switch uses.
<b>Standby Adapters</b>	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 standard switch now claims.

- 10 Click **Close** to exit the dialog box.

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



# Setting Up Networking with vSphere Distributed Switches

---

# 4

With vSphere distributed switches you can set up and configure networking in a vSphere environment.

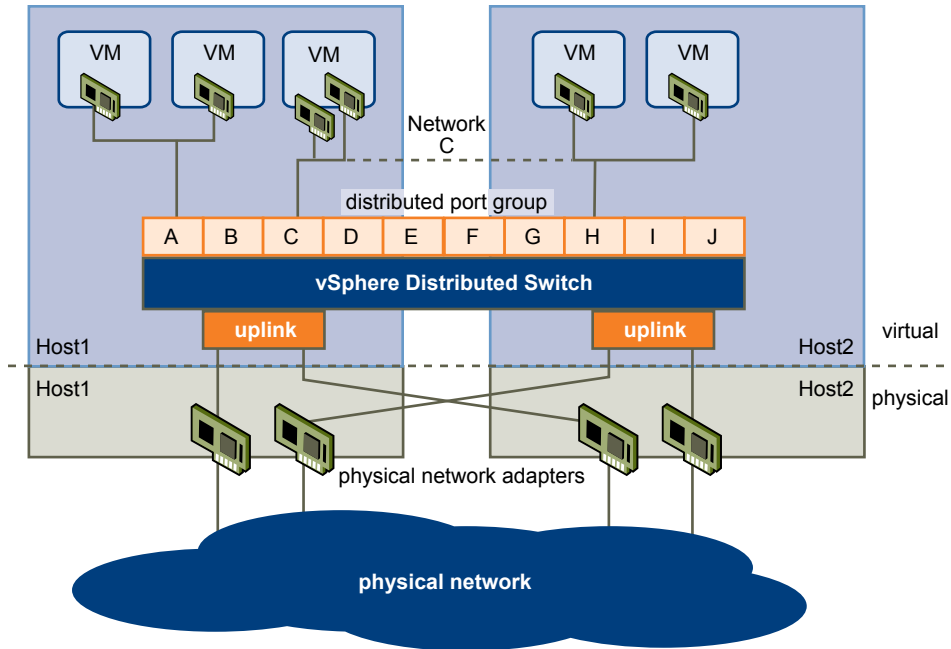
This chapter includes the following topics:

- [“vSphere Distributed Switch Architecture,”](#) on page 22
- [“Configuring a vSphere Distributed Switch,”](#) on page 22
- [“Distributed Port Groups,”](#) on page 27
- [“Working with Distributed Ports,”](#) on page 28
- [“Private VLANs,”](#) on page 29
- [“Configuring vSphere Distributed Switch Network Adapters,”](#) on page 31
- [“Configuring Virtual Machine Networking on a vSphere Distributed Switch,”](#) on page 35

## vSphere Distributed Switch Architecture

A vSphere distributed switch functions as a single switch across all associated hosts. This enables you to set network configurations that span across all member hosts, and allows virtual machines to maintain consistent network configuration as they migrate across multiple hosts.

**Figure 4-1.** vSphere Distributed Switch Network



Like a vSphere standard switch, each vSphere distributed switch is a network hub that virtual machines can use. A 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 distributed switch can also have one or more distributed port groups assigned to it. Distributed port groups group multiple ports under a common configuration and provide a stable anchor point for virtual machines connecting to labeled networks. Each distributed port 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.

Network resource pools allow you to manage network traffic by type of network traffic.

In addition to vSphere distributed switches, vSphere 5 also provides support for third-party virtual switches. For information about configuring the Cisco Nexus 1000v switch, go to <http://www.cisco.com/go/1000vdocs>.

## Configuring a vSphere Distributed Switch

You can create a vSphere distributed switch on a vCenter Server datacenter. After you have created a vSphere distributed switch, you can add hosts, create distributed port groups, and edit distributed switch properties and policies.

### Add a vSphere Distributed Switch

Create a vSphere distributed switch on a vCenter Server datacenter to handle networking traffic for all associated hosts on the datacenter.

If your system has complex port group requirements, create a distributed port group rather than a default port group.

**Procedure**

- 1 In the vSphere Client, select the Networking inventory view and select the datacenter.
- 2 Select **Inventory > Datacenter > New vSphere Distributed Switch**.
- 3 Select a vSphere distributed switch version.

Option	Description
<b>vSphere Distributed Switch Version: 4.0</b>	Compatible with ESX/ESXi version 4.0 and later. Features released with later vSphere distributed switch versions are not supported.
<b>vSphere Distributed Switch Version: 4.1.0</b>	Compatible with ESX/ESXi version 4.1 and later. Features released with later vSphere distributed switch versions are not supported.
<b>vSphere Distributed Switch Version: 5.0.0</b>	Compatible with ESXi version 5.0 and later.

- 4 Click **Next**.
- 5 In the **Name** text box, type a name for the new vSphere distributed switch.
- 6 Use the arrow buttons to select the **Number of uplink ports**, and click **Next**.  
Uplink ports connect the distributed switch to physical NICs on associated hosts. The number of uplink ports is the maximum number of allowed physical connections to the distributed switch per host.
- 7 Select whether to add hosts and their physical adapters to the vSphere distributed switch now or later.  
If you select **Add now**, select the hosts and physical adapters to use by clicking the check box next to each host or adapter. You can only free physical adapters to a vSphere distributed switch during distributed switch creation.
- 8 (Optional) Set the maximum number of ports on a host.
  - a Click **View Details** for the host.
  - b Select the maximum number of ports for the host from the drop-down menu.
  - c Click **OK**.
- 9 Click **Next**.
- 10 (Optional) Select whether to **Automatically create a default port group**.  
This option creates a distributed port group with default settings.
- 11 Click **Finish**.

**What to do next**

If you chose to add hosts later, you must add hosts to the distributed switch before adding network adapters. Network adapters can be added from the host configuration page of the vSphere Client, using Manage Hosts, or by using Host Profiles.

**Add Hosts to a vSphere Distributed Switch**

You can add hosts and physical adapters to a vSphere distributed switch at the distributed switch level after it is created.

**Procedure**

- 1 Log in to the vSphere Client and select the **Networking** inventory view.
- 2 Right-click the vSphere distributed switch in the inventory pane, and select **Add Host**.
- 3 Select the hosts to add.

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

You can select physical adapters that are not being used and physical adapters that are being used.

---

**NOTE** Moving a physical adapter to a distributed switch without moving any associated virtual adapters can cause those virtual adapters to lose network connectivity.

---

- 5 For each virtual adapter, select **Destination port group** and select a port group from the drop-down menu to migrate the virtual adapter to the distributed switch or select **Do not migrate**.
- 6 (Optional) Set the maximum number of ports on a host.
  - a Click **View Details** for the host.
  - b Select the maximum number of ports for the host from the drop-down menu.
  - c Click **OK**.
- 7 Click **Next**.
- 8 (Optional) Migrate virtual machine networking to the distributed switch.
  - a Select **Migrate virtual machine networking**.
  - b For each virtual machine, select **Destination port group** and select a port group from the drop-down menu or select **Do not migrate**.
- 9 Click **Next**.
- 10 (Optional) If you need to make any changes, click **Back** to the appropriate screen.
- 11 Review the settings for the distributed switch and click **Finish**.

## Manage Hosts on a vSphere Distributed Switch

You can change the configuration for hosts and physical adapters on a vSphere distributed switch after they are added to the distributed switch.

### Procedure

- 1 Log in to the vSphere Client and select the **Networking** inventory view.
- 2 Right-click the distributed switch and select **Manage Hosts**.
- 3 Select the hosts to manage and click **Next**.
- 4 Select the physical adapters to add, deselect the physical adapters to remove, and click **Next**.
- 5 For each virtual adapter, select the **Destination port group** from the drop-down menu to migrate the virtual adapter to the distributed switch or select **Do not migrate**.
- 6 Click **Next**.
- 7 Migrate virtual machine networking to the vSphere distributed switch.
  - a Select **Migrate virtual machine networking**.
  - b For each virtual machine, select the **Destination port group** from the drop-down menu or select **Do not migrate**.
- 8 Click **Next**.
- 9 (Optional) If you need to make any changes, click **Back** to the appropriate screen.
- 10 Review the settings for the distributed switch, and click **Finish**.

## Set the Number of Ports Per Host on a vSphere Distributed Switch

Set the maximum number of ports on a host to limit the number of distributed ports that can exist on one or more hosts associated with a vSphere distributed switch.

### Procedure

- 1 Log in to the vSphere Client and select the **Hosts and Clusters** inventory view.
- 2 Select the host to modify in the inventory pane.
- 3 On the host **Configuration** tab, click **Networking**.
- 4 Select the **vSphere Distributed Switch** view.
- 5 Click **Properties** next to the vSphere distributed switch to modify.
- 6 Select the maximum number of ports from the drop-down menu, and click **OK**.

### What to do next

If you are changing the maximum number of ports for a host after the host is added to the distributed switch, you must restart the host before the new maximum takes effect.

## Edit General vSphere Distributed Switch Settings

You can edit the general settings for a vSphere distributed switch, such as the distributed switch name and the number of uplink ports on the distributed switch.

### Procedure

- 1 Log in to the vSphere Client and select the **Networking** inventory view.
- 2 Right-click the vSphere distributed switch in the inventory pane, and select **Edit Settings**.
- 3 Select **General** to edit the vSphere distributed switch settings.

Option	Description
<b>Name</b>	Type the name for the distributed switch.
<b>Number of Uplink Ports</b>	Select the number of uplink ports for the distributed switch.
<b>Notes</b>	Type any notes for the distributed switch.

- 4 (Optional) Edit uplink port names.
  - a Click **Edit uplink names**.
  - b Type new names for one or more uplink ports.
  - c Click **OK**.
- 5 Click **OK**.

## Edit Advanced vSphere Distributed Switch Settings

You can change advanced vSphere distributed switch settings such as Cisco Discovery Protocol and the maximum MTU for the vSphere distributed switch.

### Procedure

- 1 Log in to the vSphere Client and select the **Networking** inventory view.
- 2 Right-click the vSphere distributed switch in the inventory pane, and select **Edit Settings**.

- 3 Select **Advanced** to edit the following vSphere distributed switch settings.

Option	Description
<b>Maximum MTU</b>	Maximum MTU size for the vSphere distributed switch.
<b>Discovery Protocol Status</b>	Choose the status for discovery protocol on the vSphere distributed switch. <ul style="list-style-type: none"> <li>■ <b>Enabled.</b> Enabled discovery protocol for the vSphere distributed switch. <ol style="list-style-type: none"> <li>1 Select <b>Cisco Discovery Protocol</b> or <b>Link Layer Discovery Protocol</b> from the <b>Type</b> drop-down menu.</li> <li>2 Set <b>Operation</b> to <b>Listen</b>, <b>Advertise</b>, or <b>Both</b>.</li> </ol> </li> <li>■ <b>Disabled.</b></li> </ul>
<b>Admin Contact Info</b>	Enter the <b>Name</b> and <b>Other Details</b> for the vSphere distributed switch administrator.

- 4 Click **OK**.

## View Network Adapter Information for a vSphere Distributed Switch

View physical network adapters and uplink assignments for a vSphere distributed switch from the networking inventory view of the vSphere Client.

### Procedure

- 1 Log in to the vSphere Client and select the **Networking** inventory view.
- 2 Right-click the vSphere distributed switch in the inventory pane, and select **Edit Settings**.
- 3 On the **Network Adapters** tab, you can view network adapter and uplink assignments for associated hosts. This tab is read-only. Distributed switch network adapters must be configured at the host level.
- 4 Click **OK**.

## Upgrade a vSphere Distributed Switch to a Newer Version

A vSphere distributed switch version 4.0 or 4.1 can be upgraded to a later version, enabling the distributed switch to take advantage of features that are only available in the later version.

### Procedure

- 1 Log in to the vSphere Client and select the **Networking** inventory view.
- 2 Select the vSphere distributed switch in the inventory pane.
- 3 On the **Summary** tab, next to **Version**, select **Upgrade**.

The upgrade wizard details the features available to the upgraded distributed switch that are not available to the earlier version.

- 4 Select the vSphere Distributed Switch version to upgrade to.

Option	Description
<b>vSphere Distributed Switch Version: 4.1.0</b>	Compatible with ESX/ESXi version 4.1 and later. Features released with later vSphere distributed switch versions are not supported.
<b>vSphere Distributed Switch Version: 5.0.0</b>	Compatible with ESXi version 5.0 and later.

- 5 Click **Next**.

The upgrade wizard lists the hosts associated with this vSphere distributed switch and whether or not they are compatible with the upgraded vSphere distributed switch version. You can proceed with the upgrade only if all hosts are compatible with the new vSphere distributed switch version.

Next to each incompatible host is the reason for the incompatibility.

- 6 Click **Next**.
- 7 Verify that the upgrade information listed is correct and click **Finish**.

## Distributed Port Groups

A distributed port group specifies port configuration options for each member port on a vSphere distributed switch. Distributed port groups define how a connection is made to a network.

### Add a Distributed Port Group

Add a distributed port group to a vSphere distributed switch to create a distributed switch network for your virtual machines.

#### Procedure

- 1 Log in to the vSphere Client and select the **Networking** inventory view.
- 2 Select **Inventory > vSphere Distributed Switch > New Port Group**.
- 3 Enter a **Name** and the **Number of Ports** for your new distributed port group.
- 4 Select a VLAN Type.

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

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

### Edit General Distributed Port Group Settings

You can edit general distributed port group settings such as the distributed port group name and port group type.

#### Procedure

- 1 Log in to the vSphere Client and select the **Networking** inventory view.
- 2 Right-click the distributed port group in the inventory pane, and select **Edit Settings**.
- 3 Select **General** to edit the following distributed port group settings.

Option	Action
<b>Name</b>	Type the name for the distributed port group.
<b>Description</b>	Type a brief description of the distributed port group.

Option	Action
<b>Number of Ports</b>	Type the number of ports on the distributed port group.
<b>Port binding</b>	Choose when ports are assigned to virtual machines connected to this distributed port group. <ul style="list-style-type: none"> <li>■ Select <b>Static binding</b> to assign a port to a virtual machine when the virtual machine connects to the distributed port group. This option is not available when the vSphere Client is connected directly to ESXi.</li> <li>■ Select <b>Dynamic binding</b> to assign a port to a virtual machine the first time the virtual machine powers on after it is connected to the distributed port group. Dynamic binding is deprecated in ESXi 5.0.</li> <li>■ Select <b>Ephemeral</b> for no port binding. This option is not available when the vSphere Client is connected directly to ESXi.</li> </ul>

- 4 Click **OK**.

## Edit Advanced Distributed Port Group Settings

You can edit advanced distributed port group settings, such as override settings and reset at disconnect.

### Procedure

- 1 Log in to the vSphere Client and select the **Networking** inventory view.
- 2 Right-click the distributed port group in the inventory pane, and select **Edit Settings**.
- 3 Select **Advanced** to edit the distributed port group properties.

Option	Description
<b>Allow override of port policies</b>	Select this option to allow distributed port group policies to be overridden on a per-port level. Click <b>Edit Override Settings</b> to select which policies can be overridden at the port level.
<b>Edit Override Settings</b>	Select which policies can be overridden at the port level.
<b>Configure reset at disconnect</b>	When a distributed port is disconnected from a virtual machine, the configuration of the distributed port is reset to the distributed port group setting. Any per-port overrides are discarded.

- 4 Click **OK**.

## Working with Distributed Ports

A distributed port is a port on a vSphere distributed switch that connects to the VMkernel or to a virtual machine's network adapter.

Default distributed port configuration is determined by the distributed port group settings, but some settings for individual distributed ports can be overridden.

## Monitor Distributed Port State

vSphere can monitor distributed ports and provide information on the current state of each port and the port's runtime statistics.

### Procedure

- 1 Log in to the vSphere Client and select the **Networking** inventory view.
- 2 Select the vSphere distributed switch in the inventory pane.
- 3 On the **Ports** tab, click **Start Monitoring Port State**.

The table on the Ports tab for the distributed switch now displays runtime statistics for each distributed port, including broadcast, multicast, and unicast ingress and egress traffic and packets.

The **State** column displays the current state for each distributed port.

**Table 4-1.** Distributed Port States

State	Description
Link Up	The link for this distributed port is up.
Link Down	The link for this distributed port is down.
Blocked	This distributed port is blocked.
--	The state of this distributed port is currently unavailable.

## Configure Distributed Port Settings

You can change general distributed port settings such as the port name and description.

### Procedure

- 1 Log in to the vSphere Client and select the **Networking** inventory view.
- 2 Select the vSphere distributed switch in the inventory pane.
- 3 On the **Ports** tab, right-click the port to modify and select **Edit Settings**.
- 4 Click **General**.
- 5 Modify the port name and description.
- 6 Click **OK**.

## Private VLANs

Private VLANs are used to solve VLAN ID limitations and 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 a host and the rest of the physical network, the physical switch connected to the host needs to be private VLAN-capable and configured with the VLAN IDs being used by ESXi 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 distributed ports to use Private VLAN functionality, you must create the necessary Private VLANs on the vSphere distributed switch to which the distributed ports are connected.

## Create a Private VLAN

You can create a private VLAN for use on a vSphere distributed switch and its associated distributed ports.

### Procedure

- 1 Log in to the vSphere Client and select the **Networking** inventory view.
- 2 Right-click the vSphere distributed switch in the inventory pane, and select **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 Log in to the vSphere Client and select the **Networking** inventory view.
- 2 Right-click the vSphere distributed switch in the inventory pane, and select **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 Log in to the vSphere Client and select the **Networking** inventory view.
- 2 Right-click the vSphere distributed switch in the inventory pane, and select **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 vSphere Distributed Switch Network Adapters

The vSphere distributed switch networking view of the host configuration page displays the configuration of the host's associated vSphere distributed switches and allows you to configure the vSphere distributed switch network adapters and uplink ports.

### Managing Physical Adapters

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

#### Add an Uplink to a vSphere Distributed Switch

For each host associated with a vSphere distributed switch, you must assign at least one physical network adapter, or uplink, to the vSphere distributed switch.

##### Procedure

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

The hardware configuration page for the selected host appears.

- 2 Click the **Configuration** tab and click **Networking**.
- 3 Select the **vSphere 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 will be removed from that switch and reassigned to this vSphere distributed switch.

- 7 Click **OK**.

#### Remove an Uplink from a vSphere Distributed Switch

You can remove an uplink, or physical network adapter, from a vSphere distributed 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 **vSphere Distributed Switch** view.
- 4 Click **Manage Physical Adapters**.
- 5 Click **Remove** to remove the uplink from the vSphere distributed switch.
- 6 Click **OK**.

## Removing NICs from Active Virtual Machines

When you remove NICs from active virtual machines, you may still see the NICs you removed reported in the vCenter client.

### Remove NICs from an active virtual machine without a guest operating system installed

You cannot remove NICs from an active virtual machine if the virtual machine has no operating system installed.

The vCenter client might report that the NIC has been removed, but you will continue to see it attached to the virtual machine.

### Remove NICs from an active virtual machine with a guest operating system installed

You can remove a NIC from an active virtual machine, but it might not be reported to the vCenter client for some time. If you open **Edit Settings** for the virtual machine, you might still see the NIC that you removed listed, even when the task is complete. The **Edit Settings** dialog box for the virtual machine does not immediately display the removed NIC.

## Managing Virtual Network Adapters

Virtual network adapters handle host network services over a vSphere distributed switch.

You can configure VMkernel virtual adapters for a host through an associated vSphere distributed switch either by creating new virtual adapters or migrating existing virtual adapters.

### Create a VMkernel Network Adapter on a vSphere 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 **Hosts and Clusters** inventory view.
- 2 Select the host in the inventory pane.
- 3 On the host **Configuration** tab, click **Networking**.
- 4 Select the vSphere Distributed Switch view.
- 5 Click **Manage Virtual Adapters**.
- 6 Click **Add**.
- 7 Select **New virtual adapter**, and click **Next**.
- 8 Select **VMkernel** and click **Next**.
- 9 Choose a distributed port or distributed port group connection for the virtual adapter.

Option	Description
<b>Select a port group</b>	Choose the distributed port group for the virtual adapter to connect to from the drop-down menu.
<b>Select port</b>	Type the port ID of the distributed port for the virtual network adapter to connect to.

- 10 Select **Use this virtual adapter for vMotion** to enable this port group to advertise itself to another ESXi 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 host. If this property is not enabled for any port group, migration with vMotion to this host is not possible.

- 11 Choose whether to **Use this virtual adapter for fault tolerance logging**.
- 12 Choose whether to **Use this virtual adapter for management traffic**, and click **Next**.
- 13 Under IP Settings, specify the IP address and subnet mask.  
IPv6 cannot be used with a dependent hardware iSCSI adapter.
- 14 Click **Edit** to set the VMkernel default gateway for VMkernel services, such as vMotion, NAS, and iSCSI.
- 15 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.
- 16 On the **Routing** tab, enter gateway information for the VMkernel. A gateway is needed for connectivity to machines not on the same IP subnet as the VMkernel.  
Static IP settings is the default. Do not use routing with software iSCSI Multipathing configurations or dependent hardware iSCSI adapters.
- 17 Click **OK**, and then click **Next**.
- 18 Click **Finish**.

### Migrate an Existing Virtual Adapter to a vSphere Distributed Switch

You can migrate an existing virtual adapter from a vSphere standard switch to a vSphere distributed switch.

#### Procedure

- 1 Log in to the vSphere Client and select the **Hosts and Clusters** inventory view.
- 2 Select the host in the inventory pane.
- 3 On the host **Configuration** tab, click **Networking**.
- 4 Select the vSphere Distributed Switch view.
- 5 Click **Manage Virtual Adapters**.
- 6 Click **Add**.
- 7 Select **Migrate existing virtual network adapters** and click **Next**.
- 8 Select one or more virtual network adapters to migrate.
- 9 For each selected adapter, choose a port group from the **Select a port group** drop-down menu.
- 10 Click **Next**.
- 11 Click **Finish**.

### Migrate a Virtual Adapter to a vSphere Standard Switch

You can migrate an existing virtual adapter from a vSphere distributed switch to a vSphere standard switch.

#### Procedure

- 1 Log in to the vSphere Client and select the **Hosts and Clusters** inventory view.
- 2 Select the host in the inventory pane.
- 3 On the host **Configuration** tab, click **Networking**.
- 4 Select the vSphere Distributed Switch view.
- 5 Click **Manage Virtual Adapters**.
- 6 Select the virtual adapter to migrate, and click **Migrate**.
- 7 Select the standard switch to migrate the adapter to and click **Next**.

- 8 Enter a **Network Label** and optionally a **VLAN ID** for the virtual adapter, and click **Next**.
- 9 Click **Finish** to migrate the virtual adapter and complete the wizard.

## Edit VMkernel Configuration on a vSphere Distributed Switch

You can edit a VMkernel virtual network adapter on a vSphere distributed switch to change the IP settings, such as IP address, subnet mask, default gateway, and DNS configuration. You can also select whether the virtual adapter is used for vMotion or fault tolerance logging.

### Procedure

- 1 Log in to the vSphere Client and select the **Hosts and Clusters** inventory view.
- 2 Select the host in the inventory pane.
- 3 On the host **Configuration** tab, click **Networking**.
- 4 Select the vSphere Distributed Switch view.
- 5 Click **Manage Virtual Adapters**.
- 6 Select the VMkernel adapter to modify and click **Edit**.
- 7 Under Network Connection, select **vSphere Distributed Switch** and **Port Group** or **Port** to add this virtual adapter to.
- 8 Select **Use this virtual adapter for vMotion** to enable this port group to advertise itself to another host as the network connection that vMotion traffic should be sent through.

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.

- 9 (Optional) Select **Use this virtual adapter for fault tolerance logging**.
- 10 (Optional) Select **Use this virtual adapter for management traffic**.
- 11 Under IP Settings, specify the **IP Address** and **Subnet Mask**, or select **Obtain IP settings automatically**.
- 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 appears in the name field by default. The DNS server addresses that were specified during installation are also preselected, as is the domain.

On the **Routing** tab, a gateway is needed for connectivity to machines not on the same IP subnet as the VMkernel.

Static IP settings is the default.

- 13 Use the up and down arrows to set the MTU for the VMkernel adapter.
- 14 Click **OK**.

## View VMkernel Routing Information on a vSphere Distributed Switch

You can view IP and IPv6 routing information, such as network, prefix, and gateway, for a VMkernel network adapter on a vSphere distributed switch.

### Procedure

- 1 Log in to the vSphere Client and select the **Hosts and Clusters** inventory view.
- 2 Select the host in the inventory pane.
- 3 On the host **Configuration** tab, click **Networking**.
- 4 In the vSphere Distributed Switch view, click **Manage Virtual Adapters**.

- 5 Select the VMkernel adapter to view, and click **View Routing Table** under IP Settings or IPv6 Settings. A routing table that includes network, prefix, and gateway information for the selected VMkernel adapter appears.

### Remove a Virtual Adapter

Remove a virtual network adapter from a vSphere 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 vSphere 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 vSphere Distributed Switch

Connect virtual machines to a vSphere distributed switch either by configuring an individual virtual machine NIC or migrating groups of virtual machines from the vSphere distributed switch itself.

Connect virtual machines to vSphere distributed switches by connecting their associated virtual network adapters to distributed port groups. You can do this 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 vSphere distributed switch.

### Migrate Virtual Machines to Or from a vSphere Distributed Switch

In addition to connecting virtual machines to a distributed switch at the individual virtual machine level, you can migrate a group of virtual machines between a vSphere distributed switch network and a vSphere standard switch network.

#### Procedure

- 1 Log in to the vSphere Client and select the **Networking** inventory view.
- 2 Right-click the datacenter and select **Migrate Virtual Machine Networking**.

The Migrate Virtual Machine Networking wizard appears.

- 3 Select a **Source Network** to migrate adapters from.

Option	Description
<b>Include all virtual machine network adapters that are connected to the following network (Filter by Network)</b>	Migrates virtual machine network adapters from a particular network. Select the source network from the <b>Network</b> drop-down menu.
<b>Include all virtual machine network adapters that are connected to the following network (Filter by VDS)</b>	Migrates virtual machine network adapters from a network on a particular vSphere distributed switch. To migrate from a network, select <b>Switch</b> and <b>Network</b> from the drop-down menus.
<b>Include all virtual machine network adapters that are not connected to any network</b>	Migrates virtual machine network adapters that are not connected to any network.

- 4 Select a **Destination Network** to migrate adapters to.

Option	Description
<b>Filter by Network</b>	Migrates virtual machine network adapters to a particular network. Select the destination network from the <b>Network</b> drop-down menu.
<b>Filter by VDS</b>	Migrates virtual machine network adapters to a network on a particular vSphere Distributed Switch. To migrate to a network, select <b>Switch</b> and <b>Network</b> from the drop-down menus.

- 5 Click **Next**.
- 6 (Optional) Highlight a virtual machine or adapter to view their details.
- 7 Select the virtual machines and adapters to migrate to the destination network and click **Next**.
- 8 Verify that the source network, destination network, and number of virtual machines to migrate are correct and click **OK**.

## Connect an Individual Virtual Machine to a Distributed Port Group

Connect an individual virtual machine to a vSphere 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 distributed port group to migrate to from the **Network Label** drop-down menu, and click **OK**.

# Managing Network Resources

---

vSphere provides several different methods to help you manage your network resources.

This chapter includes the following topics:

- [“vSphere Network I/O Control,”](#) on page 37
- [“TCP Segmentation Offload and Jumbo Frames,”](#) on page 40
- [“NetQueue and Networking Performance,”](#) on page 42
- [“DirectPath I/O,”](#) on page 43

## vSphere Network I/O Control

Network resource pools determine the bandwidth that different network traffic types are given on a vSphere distributed switch.

When network I/O control is enabled, distributed switch traffic is divided into the following predefined network resource pools: Fault Tolerance traffic, iSCSI traffic, vMotion traffic, management traffic, vSphere Replication (VR) traffic, NFS traffic, and virtual machine traffic.

You can also create custom network resource pools for virtual machine traffic. You can control the bandwidth each network resource pool is given by setting the physical adapter shares and host limit for each network resource pool.

The physical adapter shares assigned to a network resource pool determine the share of the total available bandwidth guaranteed to the traffic associated with that network resource pool. The share of transmit bandwidth available to a network resource pool is determined by the network resource pool's shares and what other network resource pools are actively transmitting. For example, if you set your FT traffic and iSCSI traffic resource pools to 100 shares, while each of the other resource pools is set to 50 shares, the FT traffic and iSCSI traffic resource pools each receive 25% of the available bandwidth. The remaining resource pools each receive 12.5% of the available bandwidth. These reservations apply only when the physical adapter is saturated.

---

**NOTE** The iSCSI traffic resource pool shares do not apply to iSCSI traffic on a dependent hardware iSCSI adapter.

---

The host limit of a network resource pool is the upper limit of bandwidth that the network resource pool can use.

Assigning a QoS priority tag to a network resource pool applies an 802.1p tag to all outgoing packets associated with that network resource pool.

- [Enable Network I/O Control on a vSphere Distributed Switch](#) on page 38  
Enable network resource management to use network resource pools to prioritize network traffic by type.

- [Create a Network Resource Pool](#) on page 38  
Create user-defined network resource pools for customized network resource management.
- [Add or Remove Distributed Port Groups from a Network Resource Pool](#) on page 39  
Add a distributed port group to a user-defined network resource pool to include in the network resource pool all virtual machine network traffic from that distributed port group.
- [Edit Network Resource Pool Settings](#) on page 39  
You can change network resource pool settings such as allocated shares and limits for each network resource pool to change the priority network traffic for that network resource pool is given.
- [Delete a Network Resource Pool](#) on page 40  
You can delete user-defined network resource pools that are no longer in use.

## Enable Network I/O Control on a vSphere Distributed Switch

Enable network resource management to use network resource pools to prioritize network traffic by type.

### Prerequisites

Verify that your datacenter has at least one vSphere distributed switch version 4.1.0 or later.

### Procedure

- 1 Log in to the vSphere Client and select the **Networking** inventory view.
- 2 Select the vSphere distributed switch in the inventory pane.
- 3 On the **Resource Allocation** tab, click **Properties**.
- 4 Select **Enable Network I/O Control on this vSphere distributed switch**, and click **OK**.

## Create a Network Resource Pool

Create user-defined network resource pools for customized network resource management.

User-defined network resource pools are available only on vSphere distributed switches version 5.0.0 or later.

### Procedure

- 1 Log in to the vSphere Client and select the **Networking** inventory view.
- 2 Select the vSphere distributed switch in the inventory pane.
- 3 On the **Resource Allocation** tab, click **New Network Resource Pool**.
- 4 Type a **Name** for the network resource pool.
- 5 (Optional) Type a **Description** for the network resource pool.
- 6 Select the **Physical adapter shares** for the network resource pool.

Option	Description
<b>Custom</b>	Type a specific number of shares, from 1 to 100, for this network resource pool.
<b>High</b>	Sets the shares for this resource pool to 100.
<b>Normal</b>	Sets the shares for this resource pool to 50.
<b>Low</b>	Sets the shares for this resource pool to 25.

- 7 Set the **Host limit** for the network resource pool in megabits per second or select **Unlimited**.
- 8 (Optional) Select the **QoS priority tag** for the network resource pool.

- 9 Click **OK**.

The new resource pool appears on the **Resource Allocation** tab under User-defined network resource pools.

### What to do next

Add one or more distributed port groups to the network resource pool.

## Add or Remove Distributed Port Groups from a Network Resource Pool

Add a distributed port group to a user-defined network resource pool to include in the network resource pool all virtual machine network traffic from that distributed port group.

### Prerequisites

Create one or more network resource pools on the vSphere distributed switch.

### Procedure

- 1 Log in to the vSphere Client and select the **Networking** inventory view.
- 2 Select the vSphere distributed switch in the inventory pane.
- 3 On the **Resource Allocation** tab, click **Manage Port Groups**.
- 4 (Optional) Select the user-defined network resource pool to associate with a single distributed port group from the Network resource pool drop-down menu or select **None** to remove that distributed port group from a user-defined resource pool.
- 5 (Optional) Select the user-defined network resource pool to associate with multiple distributed port groups.
  - a Hold **Ctrl** to select multiple distributed port groups to modify, and click **Assign multiple**.
  - b Select the user-defined network resource pool to associate with the distributed port groups from the Network Resource Pool drop-down menu, or select **None** to remove the distributed port groups from all user-defined resource pools.
- 6 Click **OK**.

## Edit Network Resource Pool Settings

You can change network resource pool settings such as allocated shares and limits for each network resource pool to change the priority network traffic for that network resource pool is given.

### Procedure

- 1 Log in to the vSphere Client and select the **Networking** inventory view.
- 2 Select the vSphere distributed switch in the inventory pane.
- 3 On the **Resource Allocation** tab, right-click the network resource pool to edit, and select **Edit Settings**.
- 4 Select the **Physical adapter shares** for the network resource pool.

Option	Description
<b>Custom</b>	Enter a specific number of shares, from 1 to 100, for this network resource pool.
<b>High</b>	Sets the shares for this resource pool to 100.
<b>Normal</b>	Sets the shares for this resource pool to 50.
<b>Low</b>	Sets the shares for this resource pool to 25.

- 5 Set the **Host limit** for the network resource pool in megabits per second or select **Unlimited**.

- 6 (Optional) Select the **QoS priority tag** from the drop-down menu.

The QoS priority tag specifies an IEEE 802.1p tag, allowing quality of service at the media access control level

- 7 Click **OK**.

## Delete a Network Resource Pool

You can delete user-defined network resource pools that are no longer in use.

### Prerequisites

Remove all distributed port groups from the network resource pool.

### Procedure

- 1 Log in to the vSphere Client and select the **Networking** inventory view.
- 2 Select the vSphere distributed switch in the inventory pane.
- 3 On the **Resource Allocation** tab, right-click the user-defined network resource pool to delete, and select **Delete**.
- 4 Click **Yes**.

## TCP Segmentation Offload and Jumbo Frames

You enable jumbo frames on a vSphere distributed switch or vSphere standard switch by changing the maximum transmission units (MTU). 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 replacement 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 that run 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.

## Enabling Jumbo Frames

Jumbo frames allow ESXi 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. Before enabling Jumbo frames, check with your hardware vendor to ensure that your physical network adapter supports jumbo frames.

### Enable Jumbo Frames for a VMkernel Interface on a vSphere Standard Switch

Jumbo frames reduce the CPU load caused by transferring data. Enable jumbo frames on a VMkernel network interface by changing the maximum transmission units (MTU) of the VMkernel interface.

#### Procedure

- 1 Log in to the vSphere Client and select the **Hosts and Clusters** inventory view.
- 2 On the host **Configuration** tab, click **Networking**.
- 3 Click **Properties** for the vSphere standard switch associated with the VMkernel to modify.
- 4 On the **Ports** tab, select the VMkernel interface and click **Edit**.
- 5 Set the **MTU** to 9000, and click **OK**.

### Enable Jumbo Frames on a vSphere Distributed Switch

Enable a vSphere distributed switch for jumbo frames by changing the MTU size for that distributed switch.

#### Procedure

- 1 Log in to the vSphere Client and select the **Networking** inventory view.
- 2 Right-click the vSphere distributed switch in the inventory pane, and select **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 vSphere 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 standard switch or distributed switch 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.

## NetQueue and Networking Performance

NetQueue takes advantage of the ability of some network adapters to deliver network traffic to the system in multiple receive queues that can be processed separately, allowing processing to be scaled to multiple CPUs, improving receive-side networking performance.

### Enable NetQueue on a Host

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

#### Prerequisites

Familiarize yourself with the information on configuring NIC drivers in *Getting Started with vSphere Command-Line Interfaces*.

#### Procedure

- 1 In the VMware vSphere CLI, use the command `vicfg-advcfg --set true VMkernel.Boot.netNetQueueEnable`.
- 2 Use the VMware vSphere CLI to configure the NIC driver to use NetQueue.
- 3 Reboot the host.

### Disable NetQueue on a Host

NetQueue is enabled by default.

#### Prerequisites

Familiarize yourself with the information on configuring NIC drivers in *Getting Started with vSphere Command-Line Interfaces*.

**Procedure**

- 1 In the VMware vSphere CLI, use the command `vicfg-advcfg --set false VMkernel.Boot.netNetQueueEnable`.
- 2 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`.
- 3 Reboot the host.

**DirectPath I/O**

DirectPath I/O allows virtual machine access to physical PCI functions on platforms with an I/O Memory Management Unit.

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

- 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)
- Snapshots

The following features are only available for virtual machines configured with DirectPath I/O on Cisco Unified Computing Systems (UCS) through Cisco Virtual Machine Fabric Extender (VM-FEX) distributed switches.

- vMotion
- Hot adding and removing of virtual devices
- Suspend and resume
- High availability
- DRS
- Snapshots

See Cisco VM-FEX documentation for details on supported switches and switch configuration information.

- [Configure Passthrough Devices on a Host](#) on page 43  
You can configure passthrough networking devices on a host.
- [Configure a PCI Device on a Virtual Machine](#) on page 44  
You can configure a passthrough PCI device on a virtual machine.
- [Enable DirectPath I/O with vMotion on a Virtual Machine](#) on page 44  
You can enable DirectPath I/O with vMotion for virtual machines in a datacenter on a Cisco UCS system that has at least one supported Cisco UCS Virtual Machine Fabric Extender (VM-FEX) distributed switch.

**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 DirectPath device to a virtual machine sets memory reservation to the memory size of the virtual machine.

## Enable DirectPath I/O with vMotion on a Virtual Machine

You can enable DirectPath I/O with vMotion for virtual machines in a datacenter on a Cisco UCS system that has at least one supported Cisco UCS Virtual Machine Fabric Extender (VM-FEX) distributed switch.

### Prerequisites

- Enable high performance network I/O on at least one Cisco UCS port profile on a supported Cisco VM-FEX distributed switch. For supported switches and switch configuration, see Cisco's documentation at <http://www.cisco.com/go/unifiedcomputing/b-series-doc>.
- Power off the virtual machine.

### Procedure

- 1 Log in to the vSphere Client and select the VMs and Templates inventory view.
- 2 Right-click the virtual machine to modify and click **Edit Settings**.
- 3 On the **Resources** tab, select **Memory**.
- 4 Select **Unlimited**.
- 5 On the **Hardware** tab, select the network adapter to configure as a passthrough device.
- 6 Select a port profile with high performance enabled from the network label drop-down menu, and click **OK**.
- 7 Power on the virtual machine.

After the virtual machine is powered on, DirectPath I/O appears as Active on the **Hardware** tab of the virtual machine properties dialog box.

# Networking Policies

---

Policies set at the standard switch or distributed port group level apply to all of the port groups on the standard switch or to ports in the distributed port group. The exceptions are the configuration options that are overridden at the standard port group or distributed port level.

This chapter includes the following topics:

- [“Load Balancing and Failover Policy,”](#) on page 45
- [“VLAN Policy,”](#) on page 52
- [“Security Policy,”](#) on page 52
- [“Traffic Shaping Policy,”](#) on page 56
- [“Resource Allocation Policy,”](#) on page 59
- [“Monitoring Policy,”](#) on page 60
- [“Port Blocking Policies,”](#) on page 61
- [“Manage Policies for Multiple Port Groups on a vSphere Distributed Switch,”](#) on page 62

## 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 associated with a switch or port group.

---

**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 Failover and Load Balancing Policy for a vSphere Standard Switch

Use Load Balancing and Failover policies to determine how network traffic is distributed between adapters and how to reroute traffic in the event of an adapter failure.

The Failover and Load Balancing policies include the following parameters:

- Load Balancing policy: The Load Balancing policy determines how outgoing traffic is distributed among the network adapters assigned to a standard switch. Incoming traffic is controlled by the Load Balancing policy on the physical switch.
- Failover Detection: Link Status/Beacon Probing
- Network Adapter Order (Active/Standby)

In some cases, you might lose standard switch connectivity when a failover or fallback event occurs. This causes the MAC addresses used by virtual machines associated with that standard switch to appear on a different switch port than they previously did. To avoid this problem, put your physical switch in portfast or portfast trunk mode.

### Procedure

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

The hardware configuration page for this server appears.

- 2 Click the **Configuration** tab and click **Networking**.
- 3 Select a standard switch and click **Edit**.
- 4 Click the **Ports** tab.
- 5 To edit the **Failover and Load Balancing** values, select the standard switch 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 standard switch and its port group unless you specify otherwise.

- 7 In the **Load Balancing** list, select an option for how to select an uplink.

Option	Description
<b>Route based on the originating port ID</b>	Select an uplink based on the virtual port where the traffic entered the standard switch.
<b>Route based on ip hash</b>	Select 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.
<b>Route based on source MAC hash</b>	Select an uplink based on a hash of the source Ethernet.
<b>Use explicit failover order</b>	Always use the highest order uplink from the list of Active adapters that passes failover detection criteria.

- 8 In the Network failover detection list, select the option to use for failover detection.

Option	Description
<b>Link Status only</b>	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 misconfigured to the wrong VLAN or cable pulls on the other side of a physical switch.
<b>Beacon Probing</b>	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 option detects many of the failures mentioned above that are not detected by link status alone. <b>NOTE</b> Do not use beacon probing with IP-hash load balancing.

- 9 Select **Yes** or **No** to notify switches in the case of failover.

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

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

- 10 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**, the adapter is returned to active duty immediately on 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 active adapter fails, requiring its replacement.

- 11 Set **Failover Order** to specify how to distribute the work load for adapters.

To use some adapters but reserve others for emergencies, you can set this condition using the drop-down menu to place them into groups.

Option	Description
<b>Active Adapters</b>	Continue to use the adapter when the network adapter connectivity is available and active.
<b>Standby Adapters</b>	Use this adapter if one of the active adapter's connectivity is unavailable.
<b>Unused Adapters</b>	Do not use this adapter.

If you are using iSCSI Multipathing, your VMkernel interface must be configured to have one active adapter and no standby adapters. See the *vSphere Storage* documentation.

**NOTE** When using IP-hash load balancing, do not configure standby uplinks.

## Edit the Failover and Load Balancing Policy on a Standard Port Group

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 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 port group, select the port group 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 standard switch and its port group unless you specify otherwise.

- 7 Specify the settings in the Policy Exceptions group.

Option	Description
<b>Load Balancing</b>	<p>Specify how to choose an uplink.</p> <ul style="list-style-type: none"> <li>■ <b>Route based on the originating port ID.</b> Choose an uplink based on the virtual port where the traffic entered the virtual switch.</li> <li>■ <b>Route based on ip hash.</b> 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.</li> <li>■ <b>Route based on source MAC hash.</b> Choose an uplink based on a hash of the source Ethernet.</li> <li>■ <b>Use explicit failover order.</b> Always use the highest order uplink from the list of Active adapters which passes failover detection criteria.</li> </ul> <p><b>NOTE</b> IP-based teaming requires that the physical switch be configured with etherchannel. For all other options, etherchannel should be disabled.</p>
<b>Network Failover Detection</b>	<p>Specify the method to use for failover detection.</p> <ul style="list-style-type: none"> <li>■ <b>Link Status only.</b> 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.</li> <li>■ <b>Beacon Probing.</b> 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.</li> </ul>
<b>Notify Switches</b>	<p>Select <b>Yes</b> or <b>No</b> to notify switches in the case of failover.</p> <p>If you select <b>Yes</b>, whenever a virtual NIC is connected to the standard switch 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.</p> <p><b>NOTE</b> 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.</p>
<b>Failback</b>	<p>Select <b>Yes</b> or <b>No</b> to disable or enable failback.</p> <p>This option determines how a physical adapter is returned to active duty after recovering from a failure. If failback is set to <b>Yes</b> (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 <b>No</b>, a failed adapter is left inactive even after recovery until another currently active adapter fails, requiring its replacement.</p>
<b>Failover Order</b>	<p>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:</p> <ul style="list-style-type: none"> <li>■ <b>Active Uplinks.</b> Continue to use the uplink when the network adapter connectivity is up and active.</li> <li>■ <b>Standby Uplinks.</b> Use this uplink if one of the active adapter's connectivity is down.</li> <li>■ <b>Unused Uplinks.</b> Do not use this uplink.</li> </ul>

- 8 Click **OK**.

## Edit the Teaming and Failover Policy on a Distributed Port 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 Log in to the vSphere Client and select the **Networking** inventory view.
- 2 Right-click the distributed port group in the inventory pane, and select **Edit Settings**.
- 3 Select **Policies**.
- 4 In the Teaming and Failover group specify the following.

Option	Description
<b>Load Balancing</b>	<p>Specify how to choose an uplink.</p> <ul style="list-style-type: none"> <li>■ <b>Route based on the originating virtual port</b> — Choose an uplink based on the virtual port where the traffic entered the distributed switch.</li> <li>■ <b>Route based on ip hash</b> — 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.</li> <li>■ <b>Route based on source MAC hash</b> — Choose an uplink based on a hash of the source Ethernet.</li> <li>■ <b>Route based on physical NIC load</b> — Choose an uplink based on the current loads of physical NICs.</li> <li>■ <b>Use explicit failover order</b> — Always use the highest order uplink from the list of Active adapters which passes failover detection criteria.</li> </ul> <p><b>NOTE</b> IP-based teaming requires that the physical switch be configured with etherchannel. For all other options, etherchannel should be disabled.</p>
<b>Network Failover Detection</b>	<p>Specify the method to use for failover detection.</p> <ul style="list-style-type: none"> <li>■ <b>Link Status only</b> — 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.</li> <li>■ <b>Beacon Probing</b> — 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.</li> </ul> <p><b>NOTE</b> Do not use beacon probing with IP-hash load balancing.</p>
<b>Notify Switches</b>	<p>Select <b>Yes</b> or <b>No</b> to notify switches in the case of failover.</p> <p>If you select <b>Yes</b>, whenever a virtual NIC is connected to the distributed switch 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.</p> <p><b>NOTE</b> 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.</p>

Option	Description
<b>Failback</b>	<p>Select <b>Yes</b> or <b>No</b> to disable or enable failback.</p> <p>This option determines how a physical adapter is returned to active duty after recovering from a failure. If failback is set to <b>Yes</b> (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 <b>No</b>, a failed adapter is left inactive even after recovery until another currently active adapter fails, requiring its replacement.</p>
<b>Failover Order</b>	<p>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:</p> <ul style="list-style-type: none"> <li>■ <b>Active Uplinks</b> — Continue to use the uplink when the network adapter connectivity is up and active.</li> <li>■ <b>Standby Uplinks</b>— Use this uplink if one of the active adapter’s connectivity is down.</li> <li>■ <b>Unused Uplinks</b>— Do not use this uplink.</li> </ul> <p><b>NOTE</b> When using IP-hash load balancing, do not configure standby uplinks.</p>

- 5 Click **OK**.

## Edit Distributed Port 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.

### Procedure

- 1 Log in to the vSphere Client and select the **Networking** inventory view.
- 2 Select the vSphere distributed switch in the inventory pane.
- 3 On the **Ports** tab, right-click the port to modify and select **Edit Settings**.
- 4 Click **Policies** to view and modify port networking policies.

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

Option	Description
<b>Load Balancing</b>	<p>Specify how to choose an uplink.</p> <ul style="list-style-type: none"> <li>■ <b>Route based on the originating virtual port</b> — Choose an uplink based on the virtual port where the traffic entered the vSphere distributed switch.</li> <li>■ <b>Route based on ip hash</b> — 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.</li> <li>■ <b>Route based on source MAC hash</b> — Choose an uplink based on a hash of the source Ethernet.</li> <li>■ <b>Route based on physical NIC load</b> — Choose an uplink based on the current loads of physical NICs.</li> <li>■ <b>Use explicit failover order</b> — Always use the highest order uplink from the list of Active adapters which passes failover detection criteria.</li> </ul> <p><b>NOTE</b> IP-based teaming requires that the physical switch be configured with etherchannel. For all other options, etherchannel should be disabled.</p>
<b>Network Failover Detection</b>	<p>Specify the method to use for failover detection.</p> <ul style="list-style-type: none"> <li>■ <b>Link Status only</b> — 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.</li> <li>■ <b>Beacon Probing</b> — 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.</li> </ul> <p><b>NOTE</b> Do not choose beacon probing with IP-hash load balancing.</p>
<b>Notify Switches</b>	<p>Select <b>Yes</b> or <b>No</b> to notify switches in the case of failover.</p> <p>If you select <b>Yes</b>, whenever a virtual NIC is connected to the vSphere distributed switch 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.</p> <p><b>NOTE</b> 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.</p>
<b>Failback</b>	<p>Select <b>Yes</b> or <b>No</b> to disable or enable failback.</p> <p>This option determines how a physical adapter is returned to active duty after recovering from a failure. If failback is set to <b>Yes</b> (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 <b>No</b>, a failed adapter is left inactive even after recovery until another currently active adapter fails, requiring its replacement.</p>
<b>Failover Order</b>	<p>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:</p> <ul style="list-style-type: none"> <li>■ <b>Active Uplinks</b> — Continue to use the uplink when the network adapter connectivity is up and active.</li> <li>■ <b>Standby Uplinks</b> — Use this uplink if one of the active adapter's connectivity is down.</li> </ul> <p><b>NOTE</b> When using IP-hash load balancing, do not configure standby uplinks.</p> <ul style="list-style-type: none"> <li>■ <b>Unused Uplinks</b> — Do not use this uplink.</li> </ul>

- 6 Click **OK**.

## VLAN Policy

The VLAN policy allows virtual networks to join physical VLANs.

### Edit the VLAN Policy on a Distributed Port Group

The VLAN policy allows virtual networks to join physical VLANs.

#### Procedure

- 1 Log in to the vSphere Client and select the **Networking** inventory view.
- 2 Right-click the distributed port group in the inventory pane, and select **Edit Settings**.
- 3 Select **Policies**.
- 4 Select the **VLAN Type** to use.

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

- 5 Click **OK**.

### Edit Distributed Port or Uplink Port VLAN Policies

The VLAN policy allows virtual networks to join physical VLANs.

#### Procedure

- 1 Log in to the vSphere Client and select the **Networking** inventory view.
- 2 Select the vSphere distributed switch in the inventory pane.
- 3 On the **Ports** tab, right-click the port to modify and select **Edit Settings**.
- 4 Click **Policies**.
- 5 Select the **VLAN Type** to use.

Option	Action
<b>None</b>	Do not use VLAN.
<b>VLAN</b>	In the <b>VLAN ID</b> field, enter a number between 1 and 4094.
<b>VLAN Trunking</b>	Enter one or more <b>VLAN trunk range</b> .
<b>Private VLAN</b>	Select an available private VLAN to use.

- 6 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 Security Policy for a vSphere Standard Switch

You can edit Layer 2 security policies, such as MAC address changes and forged transmits, for a vSphere standard switch.

Layer 2 is the data link layer. The three elements of the Layer 2 Security policy are promiscuous mode, MAC address changes, and forged transmits. In non-promiscuous mode, a guest adapter listens to traffic only on its own MAC address. In promiscuous mode, it can listen to all the packets. By default, guest adapters are set to non-promiscuous mode.

You can override the switch-level settings for individual standard port groups by editing the settings for the port group.

For more information about security, see the *vSphere Security* documentation.

### Procedure

- 1 Log in to the vSphere Client and select the server from the inventory panel.
- 2 Click the **Configuration** tab and click **Networking**.
- 3 Click **Properties** for the standard switch whose Layer 2 Security policy you want to edit.
- 4 In the Properties dialog box for the standard switch, click the **Ports** tab.
- 5 Select the standard switch item and click **Edit**.
- 6 Click the **Security** tab.
- 7 In the Policy Exceptions pane, select whether to reject or accept the Layer 2 Security policy exceptions.

Option	Description
<b>Promiscuous Mode</b>	<ul style="list-style-type: none"> <li>■ <b>Reject</b> — Placing a guest adapter in promiscuous mode has no effect on which frames are received by the adapter.</li> <li>■ <b>Accept</b> — Placing a guest adapter in promiscuous mode causes it to detect all frames passed on the vSphere standard switch that are allowed under the VLAN policy for the port group that the adapter is connected to.</li> </ul>
<b>MAC Address Changes</b>	<ul style="list-style-type: none"> <li>■ <b>Reject</b> — If you set the <b>MAC Address Changes</b> to <b>Reject</b> and the guest operating system changes the MAC address of the adapter to anything other than what is in the <code>.vmx</code> configuration file, all inbound frames are dropped.  If the Guest OS changes the MAC address back to match the MAC address in the <code>.vmx</code> configuration file, inbound frames are passed again.</li> <li>■ <b>Accept</b> — Changing the MAC address from the Guest OS has the intended effect: frames to the new MAC address are received.</li> </ul>
<b>Forged Transmits</b>	<ul style="list-style-type: none"> <li>■ <b>Reject</b> — Any outbound frame with a source MAC address that is different from the one currently set on the adapter are dropped.</li> <li>■ <b>Accept</b> — No filtering is performed and all outbound frames are passed.</li> </ul>

- 8 Click **OK**.

## Edit the Layer 2 Security Policy Exception for a Standard Port Group

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

### Procedure

- 1 Log in to the vSphere Client and select the **Hosts and Clusters** inventory view.
- 2 Select the host in the inventory pane.

- 3 On the host **Configuration** tab, click **Networking**.
- 4 Choose the vSphere Standard Switch view and click **Properties** for the port group to edit.
- 5 In the Properties dialog box, click the **Ports** tab.
- 6 Select the port group item and click **Edit**.
- 7 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 standard switch level.
- 8 In the Policy Exceptions pane, select whether to reject or accept the security policy exceptions.

**Table 6-1.** Policy Exceptions

Mode	Reject	Accept
Promiscuous Mode	Placing a guest adapter in promiscuous mode has no effect on which frames are received by the adapter.	Placing a guest adapter in promiscuous mode causes it to detect all frames passed on the standard switch that are allowed under the VLAN policy for the port group that the adapter is connected to.
MAC Address Changes	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.	If the MAC address from the guest OS changes, frames to the new MAC address are received.
Forged Transmits	Outbound frames with a source MAC address that is different from the one set on the adapter are dropped.	No filtering is performed, and all outbound frames are passed.

- 9 Click **OK**.

## Edit the Security Policy for a Distributed Port Group

You can set a security policy on a distributed port group to override the policy set for the distributed switch. The three elements of the Security policy are promiscuous mode, MAC address changes, and forged transmits. In nonpromiscuous mode, a guest adapter listens to traffic only on its own MAC address. In promiscuous mode, it can listen to all the packets. By default, guest adapters are set to non-promiscuous mode.

### Procedure

- 1 Log in to the vSphere Client and select the **Networking** inventory view.
- 2 Right-click the distributed port group in the inventory pane, and select **Edit Settings**.
- 3 Select **Policies**.  
By default, **Promiscuous Mode** is set to **Reject**. **MAC Address Changes** and **Forced Transmits** are set to **Accept**.

- 4 In the **Security** group, select whether to reject or accept the Security policy exceptions.

Option	Description
<b>Promiscuous Mode</b>	<ul style="list-style-type: none"> <li>■ <b>Reject</b> — Placing a guest adapter in promiscuous mode has no effect on which frames are received by the adapter.</li> <li>■ <b>Accept</b> — Placing a guest adapter in promiscuous mode causes it to detect all frames passed on the vSphere standard switch that are allowed under the VLAN policy for the port group that the adapter is connected to.</li> </ul>
<b>MAC Address Changes</b>	<ul style="list-style-type: none"> <li>■ <b>Reject</b> — If you set the <b>MAC Address Changes</b> to <b>Reject</b> and the guest operating system changes the MAC address of the adapter to anything other than what is in the .vmx configuration file, all inbound frames are dropped.  If the Guest OS changes the MAC address back to match the MAC address in the .vmx configuration file, inbound frames are passed again.</li> <li>■ <b>Accept</b> — Changing the MAC address from the Guest OS has the intended effect: frames to the new MAC address are received.</li> </ul>
<b>Forged Transmits</b>	<ul style="list-style-type: none"> <li>■ <b>Reject</b> — Any outbound frame with a source MAC address that is different from the one currently set on the adapter are dropped.</li> <li>■ <b>Accept</b> — No filtering is performed and all outbound frames are passed.</li> </ul>

- 5 Click **OK**.

## Edit Distributed Port Security Policies

The three elements of the Security policy are promiscuous mode, MAC address changes, and forged transmits.

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

### Procedure

- 1 Log in to the vSphere Client and select the **Networking** inventory view.
- 2 Right-click the vSphere distributed switch in the inventory pane, and select **Edit Settings**.
- 3 On the **Ports** tab, right-click the port to modify and select **Edit Settings**.
- 4 Click **Policies**.

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

- 5 In the **Security** group, select whether to reject or accept the Security policy exceptions.

Option	Description
<b>Promiscuous Mode</b>	<ul style="list-style-type: none"> <li>■ <b>Reject</b> — Placing a guest adapter in promiscuous mode has no effect on which frames are received by the adapter.</li> <li>■ <b>Accept</b> — Placing a guest adapter in promiscuous mode causes it to detect all frames passed on the vSphere distributed switch that are allowed under the VLAN policy for the port group that the adapter is connected to.</li> </ul>
<b>MAC Address Changes</b>	<ul style="list-style-type: none"> <li>■ <b>Reject</b> — If you set the <b>MAC Address Changes</b> to <b>Reject</b> and the guest operating system changes the MAC address of the adapter to anything other than what is in the <code>.vmx</code> configuration file, all inbound frames are dropped.  If the Guest OS changes the MAC address back to match the MAC address in the <code>.vmx</code> configuration file, inbound frames are passed again.</li> <li>■ <b>Accept</b> — Changing the MAC address from the Guest OS has the intended effect: frames to the new MAC address are received.</li> </ul>
<b>Forged Transmits</b>	<ul style="list-style-type: none"> <li>■ <b>Reject</b> — Any outbound frame with a source MAC address that is different from the one currently set on the adapter are dropped.</li> <li>■ <b>Accept</b> — No filtering is performed and all outbound frames are passed.</li> </ul>

- 6 Click **OK**.

## Traffic Shaping Policy

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

ESXi shapes outbound network traffic on standard switches and inbound and outbound traffic on distributed switches. 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.

<b>Average Bandwidth</b>	Establishes the number of bits per second to allow across a port, averaged over time. This number is the allowed average load.
<b>Peak Bandwidth</b>	Maximum number of bits per second to allow across a port when it is sending or receiving a burst of traffic. This number limits the bandwidth that a port uses when it is using its burst bonus.
<b>Burst Size</b>	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. When the port needs more bandwidth than specified by the average bandwidth, it might be allowed to temporarily transmit data at a higher speed if a burst bonus is available. This parameter limits the number of bytes that have accumulated in the burst bonus and transfers traffic at a higher speed.

## Edit the Traffic Shaping Policy for a vSphere Standard Switch

ESXi allows you to shape outbound traffic on standard switches. The traffic shaper restricts the network bandwidth available to any port, but may also be configured to 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.

<b>Average Bandwidth</b>	Establishes the number of bits per second to allow across a port, averaged over time—the allowed average load.
<b>Burst Size</b>	The maximum number of bytes to allow in a burst. If this parameter is set, a port may gain a burst bonus when it doesn’t use all its allocated bandwidth. Whenever the port needs more bandwidth than specified by <b>Average Bandwidth</b> , it may be allowed to temporarily transmit data at a higher speed if a burst bonus is available. This parameter tops the number of bytes that may be accumulated in the burst bonus and thus transferred at a higher speed.
<b>Peak Bandwidth</b>	The maximum number of bits per second to allow across a port when it is sending a burst of traffic. This tops the bandwidth used by a port whenever it is using its burst bonus. This parameter can never be smaller than the average bandwidth.

### Procedure

- 1 Log in to the vSphere Client and select the server from the inventory panel.
- 2 Click the **Configuration** tab and click **Networking**.
- 3 Select a standard switch and click **Properties**.
- 4 Click the **Ports** tab.
- 5 Select the standard switch and click **Edit**.
- 6 Click the **Traffic Shaping** tab.
- 7 Select **Enabled** from the **Status** drop-down menu to enable traffic shaping policy exceptions.

The Status policy here is applied to each virtual adapter attached to the port group, not to the standard switch as a whole. If you enable the policy exception in the **Status** field, you set limits on the amount of networking bandwidth allocation for each virtual adapter associated with this particular port group. If you disable the policy, services have a clear connection to the physical network by default.

- 8 For each traffic shaping policy, enter a bandwidth value.

## Edit the Traffic Shaping Policy for a Standard 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 **Hosts and Clusters** inventory view.
- 2 Select the host in the inventory pane.
- 3 On the host **Configuration** tab, click **Networking**.
- 4 Choose the vSphere Standard Switch view and click **Properties** for the port group to edit.
- 5 In the Properties dialog box, click the **Ports** tab.
- 6 Select the port group item and click **Edit**.

- 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
<b>Status</b>	If you enable the policy exception in the <b>Status</b> 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.
<b>Average Bandwidth</b>	A value measured over a particular period of time.
<b>Peak Bandwidth</b>	Limits the maximum bandwidth during a burst. It can never be smaller than the average bandwidth.
<b>Burst Size</b>	Specifies how large a burst can be in kilobytes (KB).

## Edit the Traffic Shaping Policy for a Distributed Port Group

ESXi allows you to shape both inbound and outbound traffic on vSphere distributed switches. The traffic shaper restricts the network bandwidth available to any port, but may also be configured to 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.

### Procedure

- Log in to the vSphere Client and select the **Networking** inventory view.
- Right-click the distributed port group in the inventory pane, and select **Edit Settings**.
- Select **Policies**.
- In the **Traffic Shaping** group, you can configure both **Ingress Traffic Shaping** and **Egress Traffic Shaping**.

When traffic shaping is disabled, the tunable features are dimmed.

**Status** — If you enable the policy exception for either **Ingress Traffic Shaping** or **Egress Traffic Shaping** 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, clear connection to the physical network by default.

- Specify network traffic parameters.

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

- Click **OK**.

## Edit Distributed Port or Uplink Port Traffic Shaping Policies

ESXi allows you to shape both inbound and outbound traffic on vSphere distributed switches. The traffic shaper restricts the network bandwidth available to any port, but may also be configured to 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.

### Procedure

- 1 Log in to the vSphere Client and select the **Networking** inventory view.
- 2 Select the vSphere distributed switch in the inventory pane.
- 3 On the **Ports** tab, right-click the port to modify and select **Edit Settings**.
- 4 Click **Policies**.
- 5 In the **Traffic Shaping** group, you can configure both **Inbound Traffic Shaping** and **Outbound Traffic Shaping**.

When traffic shaping is disabled, the tunable features are dimmed.

**Status** — If you enable the policy exception for either **Inbound Traffic Shaping** or **Outbound Traffic Shaping** 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, clear connection to the physical network by default.

- 6 Specify network traffic parameters.
  - **Average Bandwidth** establishes the number of bits per second to allow across a port, averaged over time—the allowed average load.
  - **Peak Bandwidth** is the maximum number of bits per second to allow across a port when it is sending/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 may gain a burst bonus when it doesn’t use all its allocated bandwidth. Whenever the port needs more bandwidth than specified by **Average Bandwidth**, it may be allowed to temporarily transmit data at a higher speed if a burst bonus is available. This parameter tops the number of bytes that may be accumulated in the burst bonus and thus transferred at a higher speed.
- 7 Click **OK**.

## Resource Allocation Policy

The Resource Allocation policy allows you to associate a distributed port or port group with a user-created network resource pool. This policy provides you with greater control over the bandwidth given to the port or port group.

For information about creating and configuring network resource pools, see [“vSphere Network I/O Control,”](#) on page 37.

## Edit the Resource Allocation Policy on a Distributed Port Group

Associate a distributed port group with a network resource pool to give you greater control over the bandwidth given to the distributed port group.

### Prerequisites

Enable Network I/O Control on the host and create one or more user-defined network resource pools.

### Procedure

- 1 Log in to the vSphere Client and select the **Networking** inventory view.
- 2 Right-click the distributed port group in the inventory pane, and select **Edit Settings**.
- 3 Select **Policies**.
- 4 In the Resource Allocation group, select the **Network Resource Pool** to associate the distributed port group with from the drop-down menu.
- 5 Click **OK**.

## Edit the Resource Allocation Policy on a Distributed Port

Associate a distributed port with a network resource pool to give you greater control over the bandwidth given to the port.

### Prerequisites

Enable Network I/O Control on the host and create one or more user-defined network resource pools.

### Procedure

- 1 Log in to the vSphere Client and select the **Networking** inventory view.
- 2 Select the vSphere distributed switch in the inventory pane.
- 3 On the **Ports** tab, right-click the port to modify and select **Edit Settings**.
- 4 Select **Policies**.
- 5 In the Resource Allocation group, select the **Network Resource Pool** to associate the port with from the drop-down menu.
- 6 Click **OK**.

## Monitoring Policy

The monitoring policy enables or disables NetFlow monitoring on a distributed port or port group.

NetFlow settings are configured at the vSphere distributed switch level. See [“Configure NetFlow Settings,”](#) on page 72.

## Edit the Monitoring Policy on a Distributed Port Group

With the Monitoring policy, you can enable or disable NetFlow monitoring on a distributed port group.

### Procedure

- 1 Log in to the vSphere Client and select the **Networking** inventory view.
- 2 Right-click the distributed port group in the inventory pane, and select **Edit Settings**.
- 3 Select **Policies**.

- In the Monitoring group, select the **NetFlow status**.

Option	Description
<b>Disabled</b>	NetFlow is disabled on the distributed port group.
<b>Enabled</b>	NetFlow is enabled on the distributed port group. You can configure NetFlow settings at the vSphere distributed switch level. See <a href="#">“Configure NetFlow Settings,”</a> on page 72.

- Click **OK**.

## Edit the Monitoring Policy on a Distributed Port

With the Monitoring policy, you can enable or disable NetFlow monitoring on a distributed port.

### Procedure

- Log in to the vSphere Client and select the **Networking** inventory view.
- Select the vSphere distributed switch in the inventory pane.
- On the **Ports** tab, right-click the port to modify and select **Edit Settings**.
- Select **Policies**.
- In the Monitoring group, select **NetFlow status**.

Option	Description
<b>Disabled</b>	NetFlow is disabled on the port.
<b>Enabled</b>	NetFlow is enabled on the port. You can configure NetFlow settings at the distributed switch level. See <a href="#">“Configure NetFlow Settings,”</a> on page 72.

- Click **OK**.

## Port Blocking Policies

Port blocking policies allow you to selectively block ports from sending or receiving data.

### Edit the Port Blocking Policy for a Distributed Port Group

The Miscellaneous policies dialog allows you to configure various distributed port group policies.

### Procedure

- Log in to the vSphere Client and select the **Networking** inventory view.
- Right-click the distributed port group in the inventory pane, and select **Edit Settings**.
- Select **Policies**.
- In the **Miscellaneous** group, choose whether to **Block all ports** in this distributed port group.
- Click **OK**.

### Edit Distributed Port or Uplink Port Blocking Policies

The Miscellaneous policies dialog allows you to configure distributed port or uplink port blocking policies.

### Procedure

- Log in to the vSphere Client and select the **Networking** inventory view.
- Select the vSphere distributed switch in the inventory pane.

- 3 On the **Ports** tab, right-click the port to modify and select **Edit Settings**.
- 4 Click **Policies**.
- 5 In the **Miscellaneous** group, select whether to **Block** this port.
- 6 Click **OK**.

## Manage Policies for Multiple Port Groups on a vSphere Distributed Switch

You can modify networking policies for multiple port groups on a distributed switch.

### Prerequisites

Create a vSphere distributed switch with one or more port groups.

### Procedure

- 1 Log in to the vSphere Client and select the **Networking** inventory view.
- 2 Right-click the distributed switch and select **Manage Port Groups**.
- 3 Select the policy categories to modify.

Option	Description
<b>Security</b>	Set MAC address changes, forged transmits, and promiscuous mode for the selected port groups.
<b>Traffic Shaping</b>	Set the average bandwidth, peak bandwidth, and burst size for inbound and outbound traffic on the selected port groups.
<b>VLAN</b>	Configure how the selected port groups connect to physical VLANs.
<b>Teaming and Failover</b>	Set load balancing, failover detection, switch notification, and failover order for the selected port groups.
<b>Resource Allocation</b>	Set network resource pool association for the selected port groups. This option is available for vSphere distributed switch versions 5.0.0 and later only.
<b>Monitoring</b>	Enable or disable NetFlow on the selected port groups. This option is available for vSphere distributed switch versions 5.0.0 and later only.
<b>Miscellaneous</b>	Enable or disable port blocking on the selected port groups.

- 4 Click **Next**.
- 5 Select one or more port groups to modify and click **Next**.

The policy configuration page appears. Only the policy categories you previously selected are displayed.

- 6 (Optional) In the Security group, select whether to reject or accept the Security policy exceptions.

Option	Description
<b>Promiscuous Mode</b>	<ul style="list-style-type: none"> <li>■ <b>Reject</b> — Placing a guest adapter in promiscuous mode has no effect on which frames are received by the adapter.</li> <li>■ <b>Accept</b> — Placing a guest adapter in promiscuous mode causes it to detect all frames passed on the vSphere distributed switch that are allowed under the VLAN policy for the port group that the adapter is connected to.</li> </ul>
<b>MAC Address Changes</b>	<ul style="list-style-type: none"> <li>■ <b>Reject</b> — If you set the <b>MAC Address Changes</b> to <b>Reject</b> and the guest operating system changes the MAC address of the adapter to anything other than what is in the .vmx configuration file, all inbound frames are dropped.  If the Guest OS changes the MAC address back to match the MAC address in the .vmx configuration file, inbound frames are passed again.</li> <li>■ <b>Accept</b> — Changing the MAC address from the Guest OS has the intended effect: frames to the new MAC address are received.</li> </ul>
<b>Forged Transmits</b>	<ul style="list-style-type: none"> <li>■ <b>Reject</b> — Any outbound frame with a source MAC address that is different from the one currently set on the adapter are dropped.</li> <li>■ <b>Accept</b> — No filtering is performed and all outbound frames are passed.</li> </ul>

- 7 (Optional) In the Traffic Shaping group, you can configure both **Ingress Traffic Shaping** and **Egress Traffic Shaping**.

When traffic shaping is disabled, the tunable features are dimmed.

Status — If you enable the policy exception for either **Ingress Traffic Shaping** or **Egress Traffic Shaping** in the **Status** field, you are setting limits on the amount of networking bandwidth allocated for each distributed port associated with the selected port groups. If you disable the policy, the amount of network bandwidth is not limited before it reaches the physical network .

- 8 (Optional) Specify network traffic parameters.

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

- 9 (Optional) Select the VLAN Type to use.

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

- 10 (Optional) In the Teaming and Failover group specify the following.

Option	Description
<b>Load Balancing</b>	<p>Specify how to choose an uplink.</p> <ul style="list-style-type: none"> <li>■ <b>Route based on the originating virtual port</b> — Choose an uplink based on the virtual port where the traffic entered the distributed switch.</li> <li>■ <b>Route based on ip hash</b> — 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.</li> <li>■ <b>Route based on source MAC hash</b> — Choose an uplink based on a hash of the source Ethernet.</li> <li>■ <b>Route based on physical NIC load</b> — Choose an uplink based on the current loads of physical NICs.</li> <li>■ <b>Use explicit failover order</b> — Always use the highest order uplink from the list of Active adapters which passes failover detection criteria.</li> </ul> <p><b>NOTE</b> IP-based teaming requires that the physical switch be configured with etherchannel. For all other options, etherchannel should be disabled.</p>
<b>Network Failover Detection</b>	<p>Specify the method to use for failover detection.</p> <ul style="list-style-type: none"> <li>■ <b>Link Status only</b> — 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.</li> <li>■ <b>Beacon Probing</b> — 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.</li> </ul> <p><b>NOTE</b> Do not use beacon probing with IP-hash load balancing.</p>
<b>Notify Switches</b>	<p>Select <b>Yes</b> or <b>No</b> to notify switches in the case of failover.</p> <p>If you select <b>Yes</b>, whenever a virtual NIC is connected to the distributed switch 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.</p> <p><b>NOTE</b> 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.</p>
<b>Failback</b>	<p>Select <b>Yes</b> or <b>No</b> to disable or enable failback.</p> <p>This option determines how a physical adapter is returned to active duty after recovering from a failure. If failback is set to <b>Yes</b> (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 <b>No</b>, a failed adapter is left inactive even after recovery until another currently active adapter fails, requiring its replacement.</p>
<b>Failover Order</b>	<p>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:</p> <ul style="list-style-type: none"> <li>■ <b>Active Uplinks</b> — Continue to use the uplink when the network adapter connectivity is up and active.</li> <li>■ <b>Standby Uplinks</b> — Use this uplink if one of the active adapter's connectivity is down.</li> <li>■ <b>Unused Uplinks</b> — Do not use this uplink.</li> </ul> <p><b>NOTE</b> When using IP-hash load balancing, do not configure standby uplinks.</p>

- 11 (Optional) In the Resource Allocation group, choose the **Network Resource Pool** to associate the distributed port group with from the drop-down menu.

- 12 (Optional) In the Monitoring group, choose the **NetFlow status**.

Option	Description
<b>Disabled</b>	NetFlow is disabled on the distributed port group.
<b>Enabled</b>	NetFlow is enabled on the distributed port group. NetFlow settings can be configured at the vSphere distributed switch level.

- 13 (Optional) In the **Miscellaneous** group, choose whether to **Block all ports** in this distributed port group.

- 14 Click **Next**.

All displayed policies are applied to all selected port groups, including those policies that have not been changed.

- 15 (Optional) If you need to make any changes, click **Back** to the appropriate screen.

- 16 Review the port group settings and click **Finish**.



# Advanced Networking

---

Advanced networking configuration options allow you greater control over your vSphere networking environment.

This chapter includes the following topics:

- [“Enable Internet Protocol Version 6 Support,”](#) on page 67
- [“VLAN Configuration,”](#) on page 68
- [“Working With Port Mirroring,”](#) on page 68
- [“Configure NetFlow Settings,”](#) on page 72
- [“Switch Discovery Protocol,”](#) on page 72
- [“Change the DNS and Routing Configuration,”](#) on page 74
- [“MAC Addresses,”](#) on page 74
- [“Mounting NFS Volumes,”](#) on page 76

## Enable Internet Protocol Version 6 Support

Internet Protocol version 6 (IPv6) support in ESXi provides the ability to use Virtual Infrastructure features such as NFS in an IPv6 environment. Use the Networking Properties dialog box to enable or disable IPv6 support on the host.

IPv6 is designated by the Internet Engineering Task Force as the successor to IPv4. The most obvious difference is address length. IPv6 uses 128-bit addresses rather than the 32-bit addresses used by IPv4. This increase resolves the problem of address exhaustion and eliminates the need for network address translation. Other 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.

IPv6 is disabled by default.

### Prerequisites

Required privilege: **Host.Configuration.Network Configuration**

### Procedure

- 1 From the vSphere Client Home page, click **Hosts and Clusters**.
- 2 Select the host and click the **Configuration** tab.
- 3 Click the **Networking** link under **Hardware**.
- 4 In the **vSphere Standard Switch** view, click the **Properties** link.

- 5 Select **Enable IPv6 support on this host** and click **OK**.
- 6 Reboot the host.

## VLAN Configuration

Virtual LANs (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.

Configuring ESXi with VLANs is recommended for the following reasons.

- It integrates the host into a pre-existing environment.
- It secures network traffic.
- It reduces network traffic congestion.
- iSCSI traffic requires an isolated network.

You can configure VLANs in ESXi using three methods: External Switch Tagging (EST), Virtual Switch Tagging (VST), and Virtual Guest Tagging (VGT).

With EST, all VLAN tagging of packets is performed on the physical switch. Host network adapters are connected to access ports on the physical switch. Port groups that are connected to the virtual switch must have their VLAN ID set to 0.

With VST, all VLAN tagging of packets is performed by the virtual switch before leaving the host. Host network adapters must be connected to trunk ports on the physical switch. Port groups that are connected to the virtual switch must have an appropriate VLAN ID specified.

With VGT, all VLAN tagging is performed by the virtual machine. VLAN tags are preserved between the virtual machine networking stack and external switch when frames are passed to and from virtual switches. Physical switch ports are set to trunk port.

---

**NOTE** When using VGT, you must have an 802.1Q VLAN trunking driver installed on the virtual machine.

---

## Working With Port Mirroring

Port mirroring allows you to mirror a distributed port's traffic to other distributed ports or specific physical switch ports.

### Create a Port Mirroring Session

Create a port mirroring session to mirror vSphere distributed switch traffic to specific physical switch ports.

#### Prerequisites

Create a vSphere distributed switch version 5.0.0 or later.

#### Procedure

- 1 [Specify Port Mirroring Name and Session Details](#) on page 69  
Specify the name, description, and session details for the new port mirroring session.
- 2 [Choose Port Mirroring Sources](#) on page 69  
Select sources and traffic direction for the new port mirroring session.
- 3 [Choose Port Mirroring Destinations](#) on page 69  
Select ports, or uplinks as destinations for the port mirroring session.
- 4 [Verify New Port Mirroring Settings](#) on page 70  
Verify and enable the new port mirroring session.

## Specify Port Mirroring Name and Session Details

Specify the name, description, and session details for the new port mirroring session.

### Procedure

- 1 Log in to the vSphere Client and select the **Networking** inventory view.
- 2 Right-click the vSphere distributed switch in the inventory pane, and select **Edit Settings**.
- 3 On the Port Mirroring tab, click **Add**.
- 4 Enter a **Name** and **Description** for the port mirroring session.
- 5 (Optional) Select **Allow normal IO on destination ports** to allow normal IO traffic on destination ports.  
If you do not select this option, mirrored traffic will be allowed out on destination ports, but no traffic will be allowed in.
- 6 (Optional) Select **Encapsulation VLAN** to create a VLAN ID that encapsulates all frames at the destination ports.  
If the original frames have a VLAN and **Preserve original VLAN** is not selected, the encapsulation VLAN replaces the original VLAN.
- 7 (Optional) Select **Preserve original VLAN** to keep the original VLAN in an inner tag so mirrored frames are double encapsulated.  
This option is available only if you select **Encapsulation VLAN**.
- 8 (Optional) Select **Mirrored packet length** to put a limit on the size of mirrored frames.  
If this option is selected, all mirrored frames are truncated to the specified length.
- 9 Click **Next**.

## Choose Port Mirroring Sources

Select sources and traffic direction for the new port mirroring session.

### Procedure

- 1 Choose whether to use this source for **Ingress** or **Egress** traffic, or choose **Ingress/Egress** to use this source for both types of traffic.
- 2 Type the source port IDs and click >> to add the sources to the port mirroring session.  
Separate multiple port IDs with a comma.
- 3 Click **Next**.

## Choose Port Mirroring Destinations

Select ports, or uplinks as destinations for the port mirroring session.

Port Mirroring is checked against the VLAN forwarding policy. If the VLAN of the original frames is not equal to or trunked by the destination port, the frames are not mirrored.

**Procedure**

- 1 Choose the **Source type**.

Option	Description
<b>Port</b>	Type in one or more <b>Port IDs</b> to use as a destination for the port mirroring session. Separate multiple IDs with a comma.
<b>Uplink</b>	Select one or more uplinks to use as a destination for the port mirroring session.

- 2 Click >> to add the selected destinations to the port mirroring session.
- 3 (Optional) Repeat the above steps to add multiple destinations.
- 4 Click **Next**.

**Verify New Port Mirroring Settings**

Verify and enable the new port mirroring session.

**Procedure**

- 1 Verify that the listed name and settings for the new port mirroring session are correct.
- 2 (Optional) Click **Back** to make any changes.
- 3 (Optional) Click **Enable this port mirroring session** to start the port mirroring session immediately.
- 4 Click **Finish**.

**View Port Mirroring Session Details**

View port mirroring session details, including status, sources, and destinations.

**Procedure**

- 1 Log in to the vSphere Client and select the **Networking** inventory view.
- 2 Right-click the vSphere distributed switch in the inventory pane, and select **Edit Settings**.
- 3 On the **Port Mirroring** tab, select the port mirroring session to view.  
Details for the selected port mirroring session appear under **Port Mirroring Session Details**.
- 4 (Optional) Click **Edit** to edit the details for the selected port mirroring session.
- 5 (Optional) Click **Delete** to delete the selected port mirroring session.
- 6 (Optional) Click **Add** to add a new port mirroring session.

**Edit Port Mirroring Name and Session Details**

Edit the details of a port mirroring session, including name, description, and status.

**Procedure**

- 1 Log in to the vSphere Client and select the **Networking** inventory view.
- 2 Right-click the vSphere distributed switch in the inventory pane, and select **Edit Settings**.
- 3 On the **Port Mirroring** tab, select the port mirroring session to modify and click **Edit**.
- 4 Click the **Properties** tab.
- 5 (Optional) Type a new **Name** for the port mirroring session.

- 6 (Optional) Type a new **Description** for the port mirroring session.
- 7 Select whether the port mirroring session should be **Enabled** or **Disabled**.
- 8 (Optional) Select **Allow normal IO on destination ports** to allow normal IO traffic on destination ports.  
If you do not select this option, mirrored traffic is allowed out on destination ports, but no traffic is allowed in.
- 9 (Optional) Select **Encapsulation VLAN** to create a VLAN ID that encapsulates all frames at the destination ports.  
If the original frames have a VLAN and **Preserve original VLAN** is not selected, the encapsulation VLAN replaces the original VLAN.
- 10 (Optional) Select **Preserve original VLAN** to keep the original VLAN in an inner tag so mirrored frames are double encapsulated.  
This option is available only if you select **Encapsulation VLAN**.
- 11 (Optional) Select **Mirrored packet length** to put a limit on the size of mirrored frames.  
If this option is selected, all mirrored frames are truncated to the specified length.
- 12 Click **OK**.

## Edit Port Mirroring Sources

Edit sources and traffic direction for the port mirroring session.

### Procedure

- 1 Log in to the vSphere Client and select the **Networking** inventory view.
- 2 Right-click the vSphere distributed switch in the inventory pane, and select **Edit Settings**.
- 3 On the **Port Mirroring** tab, select the port mirroring session to modify and click **Edit**.
- 4 Click the **Sources** tab.
- 5 (Optional) Select whether to use this source for **Ingress** or **Egress** traffic, or select **Ingress/Egress** to use this source for both types of traffic.
- 6 (Optional) Type one or more port IDs or ranges of port IDs to add as source for the port mirroring session, and click **>>**.  
Separate multiple IDs with commas.
- 7 (Optional) Select a source in the right-hand list and click **<<** to remove the source from the port mirroring session.
- 8 Click **OK**.

## Edit Port Mirroring Destinations

Edit the destination ports and uplinks for a port mirroring session to change where traffic for the session is mirrored.

### Procedure

- 1 Log in to the vSphere Client and select the **Networking** inventory view.
- 2 Right-click the vSphere distributed switch in the inventory pane, and select **Edit Settings**.
- 3 On the **Port Mirroring** tab, select the port mirroring session to modify and click **Edit**.
- 4 Click the **Destinations** tab.

- (Optional) Select the **Destination type** of the destination to add.

Option	Description
<b>Port</b>	Type one or more <b>Port IDs</b> to use as a destination for the port mirroring session. Separate multiple IDs with a comma.
<b>Uplink</b>	Select one or more uplinks to use as a destination for the port mirroring session.

- (Optional) Type one or more port IDs or ranges of port IDs to add as a destination for the port mirroring session and click >>.
 

Separate multiple IDs with commas.
- (Optional) Select a destination from the right-hand column and click << to remove the destination from the port mirroring session.
- Click **OK**.

## Configure NetFlow Settings

NetFlow is a network analysis tool that you can use to monitor network monitoring and virtual machine traffic.

NetFlow is available on vSphere distributed switch version 5.0.0 and later.

### Procedure

- Log in to the vSphere Client and select the **Networking** inventory view.
- Right-click the vSphere distributed switch in the inventory pane, and select **Edit Settings**.
- Navigate to the **NetFlow** tab.
- Type the **IP address** and **Port** of the NetFlow collector.
- Type the **VDS IP address**.
 

With an IP address to the vSphere distributed switch, the NetFlow collector can interact with the vSphere distributed switch as a single switch, rather than interacting with a separate, unrelated switch for each associated host.
- (Optional) Use the up and down menu arrows to set the **Active flow export timeout** and **Idle flow export timeout**.
- (Optional) Use the up and down menu arrows to set the **Sampling rate**.
 

The sampling rate determines what portion of data NetFlow collects, with the sampling rate number determining how often NetFlow collects the packets. A collector with a sampling rate of 2 collects data from every other packet. A collector with a sampling rate of 5 collects data from every fifth packet.
- (Optional) Select **Process internal flows only** to collect data only on network activity between virtual machines on the same host.
- Click **OK**.

## Switch Discovery Protocol

Switch discovery protocols allow vSphere administrators to determine which switch port is connected to a given vSphere standard switch or vSphere distributed switch.

vSphere 5.0 supports Cisco Discovery Protocol (CDP) and Link Layer Discovery Protocol (LLDP). CDP is available for vSphere standard switches and vSphere distributed switches connected to Cisco physical switches. LLDP is available for vSphere distributed switches version 5.0.0 and later.

When CDP or LLDP is enabled for a particular vSphere distributed switch or vSphere standard switch, you can view properties of the peer physical switch such as device ID, software version, and timeout from the vSphere Client.

## Enable Cisco Discovery Protocol on a vSphere Distributed Switch

Cisco Discovery Protocol (CDP) allows vSphere administrators to determine which Cisco switch port connects to a given vSphere standard switch or vSphere distributed switch. When CDP is enabled for a particular vSphere distributed switch, you can view properties of the Cisco switch (such as device ID, software version, and timeout) from the vSphere Client.

### Procedure

- 1 Log in to the vSphere Client and select the **Networking** inventory view.
- 2 Right-click the vSphere distributed switch in the inventory pane, and select **Edit Settings**.
- 3 On the **Properties** tab, select **Advanced**.
- 4 Select **Enabled** from the **Status** drop-down menu.
- 5 Select **Cisco Discovery Protocol** from the **Type** drop-down menu.
- 6 Select the CDP mode from the **Operation** drop-down menu.

Option	Description
<b>Listen</b>	ESXi detects and displays information about the associated Cisco switch port, but information about the vSphere distributed switch is not available to the Cisco switch administrator.
<b>Advertise</b>	ESXi makes information about the vSphere distributed switch available to the Cisco switch administrator, but does not detect and display information about the Cisco switch.
<b>Both</b>	ESXi detects and displays information about the associated Cisco switch and makes information about the vSphere distributed switch available to the Cisco switch administrator.

- 7 Click **OK**.

## Enable Link Layer Discovery Protocol on a vSphere Distributed Switch

With Link Layer Discovery Protocol (LLDP), vSphere administrators can determine which physical switch port connects to a given vSphere distributed switch. When LLDP is enabled for a particular distributed switch, you can view properties of the physical switch (such as chassis ID, system name and description, and device capabilities) from the vSphere Client.

LLDP is available only on vSphere distributed switch version 5.0.0 and later.

### Procedure

- 1 Log in to the vSphere Client and select the **Networking** inventory view.
- 2 Right-click the vSphere distributed switch in the inventory pane, and select **Edit Settings**.
- 3 On the **Properties** tab, select **Advanced**.
- 4 Select **Enabled** from the **Status** drop-down menu.
- 5 Select **Link Layer Discovery Protocol** from the **Type** drop-down menu.

- 6 Select the LLDP mode from the **Operation** drop-down menu.

Option	Description
<b>Listen</b>	ESXi detects and displays information about the associated physical switch port, but information about the vSphere distributed switch is not available to the switch administrator.
<b>Advertise</b>	ESXi makes information about the vSphere distributed switch available to the switch administrator, but does not detect and display information about the physical switch.
<b>Both</b>	ESXi detects and displays information about the associated physical switch and makes information about the vSphere distributed switch available to the switch administrator.

- 7 Click **OK**.

## View Switch Information on the vSphere Client

When CDP or LLDP is set to **Listen** or **Both**, you can view physical switch information from 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 **Networking**.
- 3 Click the information icon to the right of the vSphere standard switch or vSphere distributed switch to display information for that switch.

Switch information for the selected switch appears.

## 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.
- 6 Specify the domains in which to look for hosts.
- 7 On the **Routing** tab, change the default gateway information as needed.
- 8 Click **OK**.

## MAC Addresses

MAC addresses are generated for virtual network adapters that virtual machines and network services 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. By default, VMware uses the Organizationally Unique Identifier (OUI) 00:50:56 for manually generated addresses, but all unique manually generated addresses are supported.

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  $\langle number \rangle$  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 ESXi 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 ESXi

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 ESXi 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 assign static MAC addresses to a powered-down virtual machine's virtual NICs.

### 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 static MAC address, and click **OK**.

## Mounting NFS Volumes

ESXi supports VMkernel-based NFS mounts.

NFS volumes are mounted with ISO images by using VMkernel NFS. All NFS volumes mounted in this way appear as datastores in the vSphere Client.

# Networking Best Practices

---

Consider these best practices when you configure your network.

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

Put a set of virtual machines on a separate physical NIC. This separation allows for a portion of the total networking workload to be shared evenly across multiple CPUs. The isolated virtual machines can then better serve traffic from a Web client, for example
- Keep the vMotion connection on a separate network devoted to vMotion. When migration with vMotion occurs, the contents of the guest operating system's memory is transmitted over the network. You can do this either by using VLANs to segment a single physical network or by using separate physical networks (the latter is preferable).
- 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 vSphere standard switch or vSphere distributed switch for each service. If this is not possible, separate network services on a single switch 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 network adapters from a standard or distributed switch without affecting the virtual machines or the network service that is running behind that switch. If you remove all the running hardware, the virtual machines can still communicate among themselves. If you leave one network adapter 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.
- For best performance, use vmxnet3 virtual NICs.
- Every physical network adapter connected to the same vSphere standard switch or vSphere distributed switch should also be connected to the same physical network.
- Configure all VMkernel network adapters to the same MTU. When several VMkernel network adapters are connected to vSphere distributed switches but have different MTUs configured, you might experience network connectivity problems.



# Index

## A

- active adapters **19**
- active uplinks **47, 50**
- adding
  - distributed port groups **27**
  - vSphere distributed switch **22**
- adding a VMkernel network adapter **16**
- admin contact info **25**
- average bandwidth **56–59, 62**

## B

- bandwidth
  - average **56, 57**
  - peak **56, 57**
- beacon probing, standard switches **46**
- binding on host, distributed port groups **28**
- blocked ports
  - distributed port groups **61, 62**
  - distributed ports **61**
- burst size **56–59, 62**

## C

- CDP **72, 74**
- Cisco Discovery Protocol **25, 73, 74**
- Cisco switches **72**
- config reset at disconnect, distributed port groups **28**

## D

- default gateway, editing **34**
- delete resource pool, vSphere distributed switch **40**
- DirectPath I/O, vMotion **44**
- DirectPath I/O Gen. 2 **44**
- distributed port groups
  - adding **27**
  - average bandwidth **58, 62**
  - binding on host **28**
  - blocked ports **61, 62**
  - burst size **58, 62**
  - config reset at disconnect **28**
  - description **27**
  - failover order **49, 62**
  - failover policies **49, 62**
  - forged transmits **54, 62**
  - live port moving **28**

- load balancing **49, 62**
- MAC address changes **54, 62**
- miscellaneous policies **61, 62**
- name **27**
- NetFlow **60, 62**
- Network I/O Control **60, 62**
- network resource pools **60, 62**
- notify switches **49, 62**
- number of ports **27**
- override settings **28**
- peak bandwidth **58, 62**
- port group type **27**
- port name format **28**
- port policies **61, 62**
- promiscuous mode **54, 62**
- PVLAN **52, 62**
- QOS policies **52, 62**
- resource pool **39**
- security policy **54, 62**
- teaming policies **49, 62**
- traffic shaping **58, 62**
- virtual machines **36**
- VLAN policy **52, 62**
- VLAN trunking **52, 62**
- distributed ports
  - blocked ports **61**
  - blocking **61**
  - failback **50**
  - failover order **50**
  - load balancing **60**
  - monitoring **28**
  - NetFlow **61**
  - network failover detection **50**
  - Network I/O Control **60**
  - network resource pools **60**
  - notify switches **50**
  - port policies **61**
  - properties **29**
  - states **28**
  - teaming and failover policies **50**
  - traffic shaping policies **59**
  - VLAN policies **52**
- distributed switch
  - adding **22**

- adding a host to **23**
- admin contact info **25**
- Cisco Discovery Protocol **25**
- configuration **22**
- hosts **25**
- IP address **25**
- jumbo frames **41**
- maximum MTU **25**
- maximum number of ports **25**
- migrating virtual machines to or from **35**
- new resource pool **38**
- ports **25**
- resource pool settings **38**
- upgrading **26**
- virtual machines **35**
- virtual network adapter **33**
- virtual network adapters **32**

**DNS 74**

DNS configuration, vSphere distributed switch **34**

**E**

early binding port groups **27**

enhanced vmxnet **40, 41**

**F**

failback **47, 49, 50, 62**

failover **45**

failover order, distributed port groups **49, 62**

failover policies

- distributed port groups **49, 62**
- distributed ports **50**
- port group **47**
- standard switches **46**

Fault Tolerance, logging **34**

forced transmits **55**

forged transmits **53, 54, 62**

**G**

guest operating system, NICs **32**

**H**

host networking, viewing **10**

hosts, adding to a vSphere distributed switch **23**

**I**

inbound traffic shaping **59**

Internet Protocol version 6 **67**

IP address, editing **34**

IP storage port groups, creating **16, 32**

IPv6 **67**

iSCSI, networking **16, 68**

**J**

- jumbo frames
  - enabling **41**
  - virtual machines **40, 41**

**L**

- late binding port groups **27**
- Layer 2 security **52**
- Layer 2 security policy **53**
- Link Layer Discovery Protocol **72–74**
- link status, standard switches **46**
- live port moving, distributed port groups **28**
- LLDP, enable **73**
- load balancing, distributed port groups **49, 62**
- load balancing policies, standard switches **46**

**M**

- MAC address
  - configuration **76**
  - configuring **74**
  - generating **75**
  - static **76**
- MAC address changes **53, 54, 62**
- MAC addresses **55**
- maximum MTU **25**
- maximum number of ports **25**
- miscellaneous policies, distributed port groups **61, 62**
- MTU **40, 41**

**N**

- NAS, mounting **76**
- NetFlow
  - collector settings **72**
  - configure **72**
  - disable **60–62**
  - distributed port Groups **60, 62**
  - distributed ports **61**
  - enable **60–62**
  - netqueue, enable **42**
  - NetQueue, disabling **42**
  - network adapters
    - distributed switch **32**
    - viewing **11, 26**
    - vSphere distributed switch **31**
  - network failover detection **47, 50**
  - Network I/O Control **59**
  - network resource management **37**
  - network resource pools
    - distributed port groups **60, 62**
    - distributed ports **60**
  - networking
    - advanced **67**
    - introduction **9**

- performance **42**
  - security policies **55**
  - networking best practices **77**
  - networks
    - distributed ports **28**
    - resource pools **37**
    - resource settings **38–40**
  - new resource pool, distributed switch **38**
  - NFS, networking **16**
  - NIC teaming
    - definition **9**
    - standard switches **46**
  - NICs
    - adding to a vSphere distributed switch **31**
    - remove from active virtual machine **7, 32**
    - removing from a vSphere distributed switch **31, 32**
    - vCenter client **32**
  - notify switches **47, 49, 50, 62**
- O**
- outbound traffic shaping **59**
  - override settings, distributed port groups **28**
- P**
- passthrough device, add to a virtual machine **44**
  - PCI **43**
  - peak bandwidth **56–59, 62**
  - physical network adapters
    - adding to a vSphere distributed switch **31**
    - managing **31**
    - removing **31**
  - port blocking **45**
  - port configuration **18**
  - port groups
    - definition **9**
    - failback **47**
    - failover order **47**
    - Layer 2 Security **53**
    - load balancing **47**
    - network failover detection **47**
    - notify switches **47**
    - traffic shaping **57**
    - using **14**
  - port mirroring
    - create **68**
    - destinations **69–71**
    - name **69, 70**
    - packet length **69**
    - sources **69–71**
    - status **70**
    - verify settings **70**
    - VLAN **69, 70**
  - port name format, distributed port groups **28**
  - port policies, distributed port groups **61, 62**
  - ports, vSphere distributed switch **25**
  - private VLAN
    - create **29**
    - primary **30**
    - removing **30**
    - secondary **30**
  - promiscuous mode **53–55, 62**
  - properties, distributed ports **29**
  - PVLAN **52**
- Q**
- QOS policies, distributed port groups **52, 62**
- R**
- resource pool, distributed port groups **39**
  - resource pool settings
    - distributed switch **38**
    - vSphere distributed switch **39**
  - resource pools, networks **37**
  - routing **74**
- S**
- security policies, distributed ports **55**
  - security policy
    - distributed port groups **54, 62**
    - policy exceptions **53**
    - virtual switches **53**
  - standard port group **14**
  - standard switch **13**
  - standard switches
    - average bandwidth **57**
    - beacon probing **46**
    - burst size **57**
    - configuration **18**
    - failover **46**
    - forged transmits **53**
    - link status **46**
    - load balancing policies **46**
    - MAC address changes **53**
    - NIC teaming **46**
    - peak bandwidth **57**
    - port configuration **18**
    - promiscuous mode **53**
    - properties **18**
    - security policy **53**
    - traffic shaping policies **57**
    - using **13**
  - standby adapters **19**
  - standby uplinks **47, 50**
  - states, distributed ports **28**
  - subnet mask, editing **34**

**T**

- TCP Segmentation Offload **40**
- TCP/IP **16**
- teaming policies
  - distributed port groups **49, 62**
  - distributed ports **50**
  - port group **47**
- third-party switch **22**
- traffic shaping
  - distributed port groups **58, 62**
  - port groups **57**
- traffic shaping policies
  - distributed ports **59**
  - uplink ports **59**
- TSO **40**

**U**

- updated information **7**
- upgrading
  - distributed switch **26**
  - vSphere distributed switch **26**
- uplink adapters
  - adding **19**
  - adding to a vSphere distributed switch **31**
  - duplex **18**
  - managing **31**
  - removing **31**
  - speed **18**
- uplink assignments **26**
- uplink ports
  - traffic shaping policies **59**
  - VLAN policies **52**

**V**

- vCenter client **32**
- virtual adapter **33**
- Virtual LAN **68**
- virtual machine networking **10, 14, 15**
- virtual machines
  - migrating to or from a distributed switch **35**
  - migrating to or from a vSphere distributed switch **35**
  - networking **35, 36**
- virtual network adapters, removing **35**
- VLAN
  - definition **9**
  - port mirroring **69, 70**
  - private **29**
- VLAN ID
  - primary **29**
  - secondary **29**
- VLAN policies
  - distributed ports **52**
  - uplink ports **52**

- VLAN policy, distributed port groups **52, 62**
- VLAN trunking, distributed port groups **52, 62**
- VLAN Trunking **27, 52**
- VLAN Type **52**
- VMkernel
  - configuring **15**
  - definition **9**
  - gateway **17, 34**
  - jumbo frames **41**
  - networking **16**
  - prefix **17, 34**
  - routing **17, 34**
- VMkernel network adapters
  - adding **16, 32**
  - editing **34**
  - enabling vMotion **34**
  - fault tolerance logging **34**
- VMkernel networking **10**
- vMotion
  - compatibility **43**
  - definition **9**
  - DirectPath I/O **44**
  - enabling on a virtual network adapter **34**
  - networking configuration **15**
- vMotion, networking **16**
- vMotion interfaces, creating **16, 32**
- vSphere distributed switch
  - adding **22**
  - adding a host to **23**
  - adding a VMkernel network adapter **32**
  - admin contact info **25**
  - CDP **73**
  - Cisco Discovery Protocol **25, 73**
  - configuration **22**
  - delete resource pool **40**
  - editing **34**
  - hosts **25**
  - IP address **25**
  - jumbo frames **41**
  - LLDP **73**
  - manage hosts **24**
  - maximum MTU **25**
  - maximum number of ports **25**
  - migrating virtual machines to or from **35**
  - mirror **68**
  - port mirroring **68**
  - ports **25**
  - resource pool settings **39**
  - third-party **22**
  - upgrading **26**
  - virtual machines **35**

- virtual network adapter **33**
- virtual network adapters **32**
- vSphere standard switch
  - configuration **18**
  - definition **9**
  - port configuration **18**
  - properties **18**
  - teaming and failover policies **47**
  - using **13**
  - viewing **10**

