vSphere Networking

Update 1
vSphere 5.5
ESXi 5.5
vCenter Server 5.5

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

EN-001359-03

vmware®

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

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

The VMware Web site also provides the latest product updates.

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

docfeedback@vmware.com

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

Contents

About vSphere Networking

vSphere Networking provides information about configuring networking for VMware vSphere®, 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 1

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

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

Revision	Description
EN-001359-03	■ The instructions about configuring the teaming and failover policy are extended with the **Route based on physical NIC load** option. See "Edit the Teaming and Failover Policy on a Distributed Port Group in the vSphere Web Client," on page 82 and "Edit Distributed Port Teaming and Failover Policies with the vSphere Web Client," on page 84. ■ "Configure NetFlow Settings with the vSphere Web Client," on page 159 incorrectly described that the NetFlow sampling rate determines that a packet out of *x* are dropped.
EN-001359-02	■ The information about placing a host and NFS storage in different subnets is removed from "Mounting NFS Volumes," on page 161. ■ The value of the interval for beacon probing is fixed in documents within the "Teaming and Failover Policy," on page 78 chapter.
EN-001359-01	■ The instructions about changing the number of virtual switch ports on a host that runs ESXi 5.1 and earlier are removed. ■ A section about securing infrastructure traffic is added in "The VMkernel Networking Layer," on page 68. ■ The information about providing redundancy for VMkernel adapters for management traffic is updated in "The VMkernel Networking Layer," on page 68.
EN-001359-00	Initial release.

Introduction to Networking 2

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

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.

Physical Network — 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.

Virtual Network — A network of virtual machines running on a 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.

Physical Ethernet Switch — It 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.

vSphere Standard Switch — It 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.
Standard Port Group	It 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.
vSphere Distributed Switch	It acts as a single switch across all associated hosts in a datacenter to provide centralized provisioning, administration, and monitoring of virtual networks. You configure a vSphere distributed switch on the vCenter Server system and the configuration is populated across all hosts that are associated with the switch. This lets virtual machines to maintain consistent network configuration as they migrate across multiple hosts.
Host Proxy Switch	A hidden standard switch that resides on every host that is associated with a vSphere distributed switch. The host proxy switch replicates the networking configuration set on the vSphere distributed switch to the particular host.
Distributed Port	A port on a vSphere distributed switch that connects to a host's VMkernel or to a virtual machine's network adapter.
Distributed Port Group	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	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.
VLAN	VLAN 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.
VMkernel TCP/IP Networking Layer	The VMkernel networking layer provides connectivity to hosts and handles the standard infrastructure traffic of vSphere vMotion, IP storage, Fault Tolerance, and Virtual SAN.
IP Storage	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	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.

Network Services in ESXi

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

You can 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.

VLAN Configuration

Virtual LANs (VLANs) enable a single physical LAN segment to be further isolated 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 isolates and secures network traffic.
- It reduces network traffic congestion.

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 a VLAN ID between 1 and 4094.

With VGT, all VLAN tagging is done by the virtual machine. VLAN tags are preserved between the virtual machine networking stack and external switch when frames pass to and from virtual switches. Host network adapters must be connected to trunk ports on the physical switch. For a standard switch the VLAN ID of port groups with VGT must be set to 4095. For a distributed switch the VLAN trunking policy must include the range of the VLANs to which virtual machines are connected.

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

VMware ESXi Dump Collector Support

The ESXi Dump Collector sends the state of the VMkernel memory, that is, a core dump to a network server when the system encounters a critical failure.

The ESXi Dump Collector in ESXi 5.1 and later supports both vSphere Standard and Distributed Switches. The ESXi Dump Collector can also use any active uplink adapter from the team of the port group that handles the VMkernel adapter for the collector.

Changes to the IP address for the ESXi Dump Collector interface are automatically updated if the IP addresses for the configured VMkernel adapter changes. The ESXi Dump Collector also adjusts its default gateway if the gateway configuration of the VMkernel adapter changes.

If you try to delete the VMkernel network adapter used by the ESXi Dump Collector, the operation fails and a warning message appears. To delete the VMkernel network adapter, disable dump collection and delete the adapter.

There is no authentication or encryption in the file transfer session from a crashed host to the ESXi Dump Collector. You should configure the ESXi Dump Collector on a separate VLAN when possible to isolate the ESXi core dump from regular network traffic.

For information about installing and configuring the ESXi Dump Collector, see the *vSphere Installation and Setup* documentation.

Setting Up Networking with vSphere Standard Switches

3

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

This chapter includes the following topics:

vSphere Standard Switches

You can create abstracted network devices called vSphere standard switches. You use standard switches to provide network connectivity to hosts and virtual machines. A standard switch can bridge traffic internally between virtual machines in the same VLAN and link to external networks.

Standard Switch Overview

To provide network connectivity to hosts and virtual machines, you connect the physical NICs of the hosts to uplink ports on the standard switch. Virtual machines have network adapters (vNICs) that you connect to port groups on the standard switch. Every port group can use one or more physical NICs to handle their network traffic. If a port group does not have a physical NIC connected to it, virtual machines on the same port group can only communicate with each other but not with the external network.

Figure 3-1. vSphere standard switch architecture

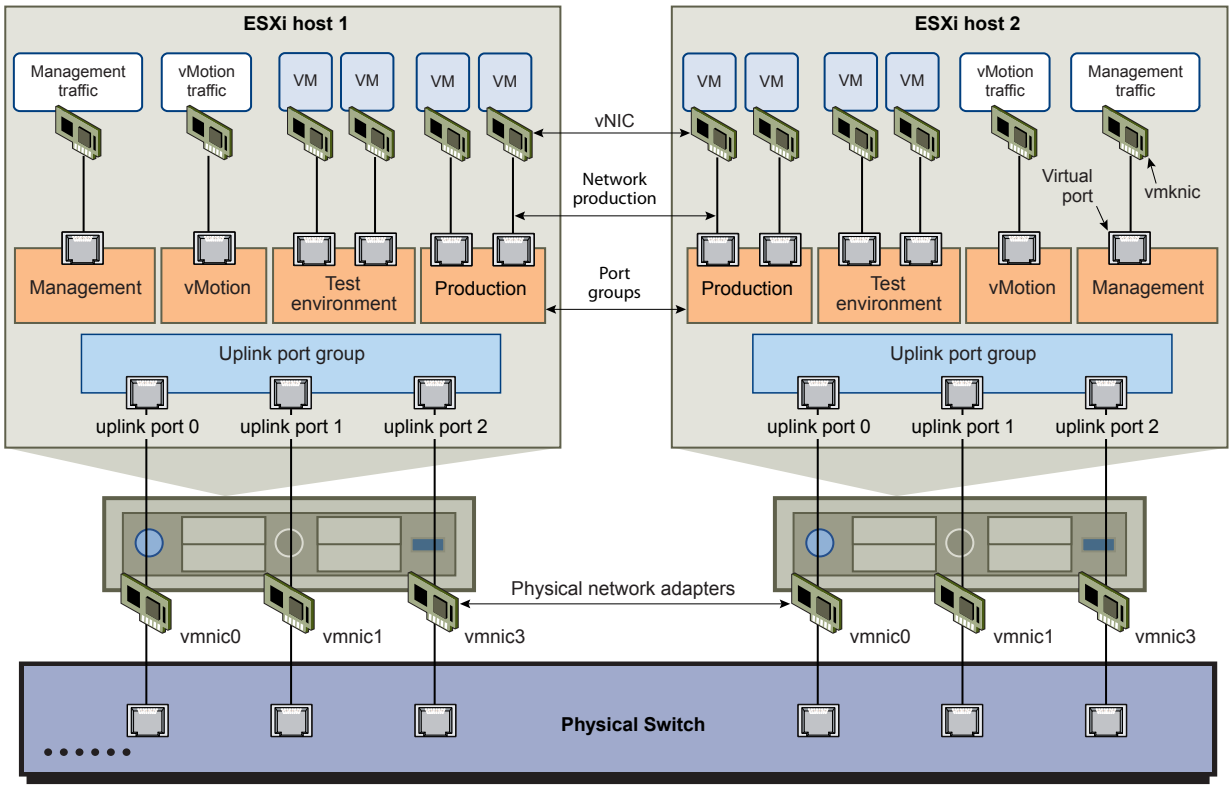

A vSphere standard switch is very similar to a physical Ethernet switch. Virtual machine network adapters and physical NICs on the host use the logical ports on the switch as each adapter uses one port. Each logical port on the standard switch is a member of a single port group. For information about maximum allowed ports and port groups, see the *Configuration Maximums* documentation.

Standard Port Groups

Each port group on a standard switch is identified by a network label, which must be unique to the current host. You can use network labels to make the networking configuration of virtual machines portable across hosts. You should give the same label to the port groups in a datacenter that use physical NICs connected to one broadcast domain on the physical network. Conversely, if two port groups are connected to physical NICs on different broadcast domains, the port groups should have distinct labels.

For example, you can create *Production* and *Test environment* port groups as virtual machine networks on the hosts that share the same broadcast domain on the physical network.

A VLAN ID, which restricts port group traffic to a logical Ethernet segment within the physical network, is optional. For port groups to receive the traffic that the same host sees, but from more than one VLAN , the VLAN ID must be set to VGT (VLAN 4095).

Number of Standard Ports

To ensure efficient use of host resources on hosts running ESXi 5.5 and later, the number of ports of standard switches are dynamically scaled up and down. A standard switch on such a host can expand up to the maximum number of ports supported on the host. The port limit is determined based on the maximum number of virtual machines that the host can handle.

Port Group Configuration for Virtual Machines

You can add or modify a virtual machine port group to set up traffic management on a set of virtual machines.

The Add Networking wizard in the vSphere Web Client guides you through the process 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 with the vSphere Web Client

Create port groups in a vSphere standard switch to provide connectivity and common network configuration for a set of virtual machines.

Procedure

1 In the vSphere Web Client, navigate to the host.

2 Right-click the host in the navigator and select **All vCenter Actions > Add Networking**.

3 In **Select connection type**, select **Virtual Machine Port Group for a Standard Switch** and click **Next**.

4 In **Select target device**, select an existing standard switch or create a new standard switch.

5 If the new port group is for an existing standard switch, navigate to the switch.

   a Click **Browse**.

   b Select a standard switch from the list and click **OK**.

   c Click **Next** and go to Step 7.

6 (Optional) In the Create a Standard Switch page, assign physical network adapters to the standard switch.

   You can create a standard switch with or without 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.

   a Click **Add adapters**.

   b Select an adapter from the **Network Adapters** list.

c Use the **Failover order group** drop-down menu to assign the adapter to Active adapters, Standby adapters, or Unused adapters, and click **OK**.

d (Optional) Use the up and down arrows in the **Assigned adapters** list to change the position of the adapter if needed.

e Click **Next**.

7 In the Connection settings page, identify traffic through the ports of the group.

a Type a **Network Label** for the port group, or accept the generated label.

b Set the **VLAN ID** to configure VLAN handling in the port group.

The VLAN ID also reflects the VLAN tagging mode in the port group.

VLAN Tagging Mode	VLAN ID	Description
External Switch Tagging (EST)	0	The virtual switch does not pass traffic associated with a VLAN.
Virtual Switch Tagging (VST)	From 1 to 4094	The virtual switch tags traffic with the entered tag.
Virtual Guest Tagging (VGT)	4095	Virtual machines handle VLANs. The virtual switch passes traffic from any VLAN.

c Click **Next**.

8 Review the port group settings in the Ready to complete page, and click **Finish**.

Click **Back** if you want to change any settings.

Edit a Standard Switch Port Group in the vSphere Web Client

By using the vSphere Web Client edit the name and VLAN ID of a standard switch port group, and override networking policies at the port group level.

Procedure

1 In the vSphere Web Client, navigate to the host.

2 On the **Manage** tab, click **Networking**, and select **Virtual switches**.

3 Select a standard switch from the list.

The topology diagram of the switch appears.

4 In the topology diagram of the switch, click the name of the port group.

5 Click **Edit** under the topology diagram title.

6 In the **Properties** section, rename the port group in the **Network Label** text field.

7 Configure VLAN tagging in the **VLAN ID** drop-down menu.

VLAN Tagging Mode	VLAN ID	Description
External Switch Tagging (EST)	0	The virtual switch does not pass traffic associated with a VLAN.
Virtual Switch Tagging (VST)	From 1 to 4094	The virtual switch tags traffic with the entered tag.
Virtual Guest Tagging (VGT)	4095	Virtual machines handle VLANs. The virtual switch passes traffic from any VLAN.

8 In the **Security** section, override the switch settings for protection against MAC address impersonation and for running virtual machines in promiscuous mode.

9 In the **Traffic Shaping** section, override at the port group level the size of average and peak bandwidth and of bursts.

10 In the **Teaming and Failover** section, override the teaming and failover settings inherited from the standard switch.

   You can configure traffic distribution and rerouting between the physical adapters associated with the port group. You can also change the order in which host physical adapters are used upon failure.

11 Click **OK**.

Remove a Port Group from a vSphere Standard Switch in the vSphere Web Client

You can remove port groups from vSphere standard switches in case you no longer need the associated labeled networks.

Prerequisites

Verify that there are no powered-on virtual machines connected to the port group that you want to remove.

Procedure

1 In the vSphere Web Client, navigate to the host.

2 On the **Manage** tab, click **Networking**, and select **Virtual switches**.

3 Select the standard switch.

4 From the topology diagram of the switch, select the port group that you want to remove by clicking its label.

5 From the toolbar in the switch topology, click the **Remove selected port group** action icon .

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 Speed of a Physical Adapter in the vSphere Web Client

A physical adapter can become a bottleneck for network traffic if the adapter speed does not match application requirements. You can change the connection speed and duplex of a physical adapter to transfer data in compliance with traffic rate.

Procedure

1 Browse to a host in the vSphere Web Client navigator.

2 Click the **Manage** tab, and select **Physical adapters** from **Networking**.

   The physical network adapters of the host appear in a table that contains details for each physical network adapter.

3 Select the physical network adapter from the list and click **Edit**.

4 Select speed and duplex mode of the physical network adapter from the drop-down menu.

5 (Optional) If the physical adapter supports SR-IOV, enable it and configure the number of virtual functions to use for virtual machine networking.

6 Click **OK**.

Add and Team Physical Adapters in a Standard Switch in the vSphere Web Client

Assign a physical adapter to a standard switch to provide connectivity to virtual machines and VMkernel adapters on the host. You can form a team of NICs to distribute traffic load and to configure failover.

NIC teaming combines multiple network connections to increase throughput and provide redundancy should a link fail. To create a team, you associate multiple physical adapters to a single vSphere standard switch.

Procedure

1 In the vSphere Web Client, navigate to the host.

2 On the **Manage** tab, click **Networking**, and select **Virtual switches**.

3 Select the standard switch you want to add a physical adapter to.

4 Click **Manage the physical network adapters**.

5 Click **Add adapters**.

6 Select one or more physical network adapters from the list and select the **Failover order group** to assign the adapters to, and click **OK**.

The selected adapters appear in the selected failover group list under the Assigned Adapters list.

7 (Optional) Use the up and down arrows to change the position of an adapter in the failover groups.

8 Click **OK** to apply the physical adapter configuration.

View the Topology Diagram of a vSphere Standard Switch in the vSphere Web Client

You can examine the structure and components of a vSphere standard switch by using its topology diagram.

The topology diagram of a standard switch provides a visual representation of the adapters and port groups connected to the switch.

From the diagram you can edit the settings of a selected port group and of a selected adapter.

Procedure

1 In the vSphere Web Client, navigate to the host.

2 On the **Manage** tab, click **Networking**, and select **Virtual switches**.

3 Select the standard switch from the list.

The diagram appears under the list of virtual switches on the host.

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

A vSphere Distributed Switch provides centralized management and monitoring of the networking configuration of all hosts that are associated with the switch. You set up a distributed switch on vCenter Server system, and its settings are propagated to all hosts that are associated with the switch.

Figure 4-1. vSphere Distributed Switch Architecture

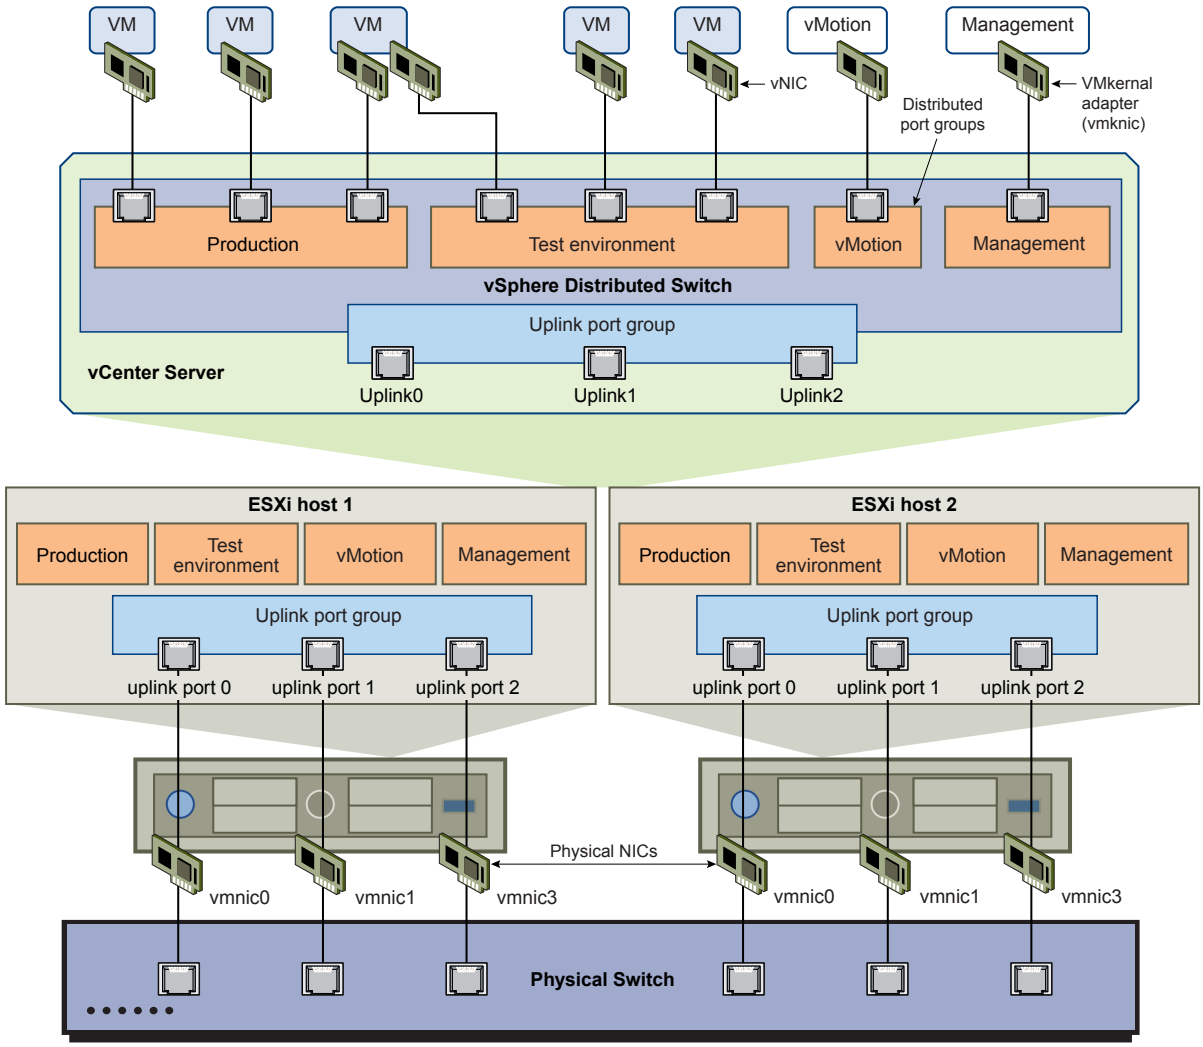

You associate a vSphere Distributed Switch with a datacenter on a vCenter Server system. The networking configuration and management for all hosts that are associated with the switch is centralized on the vCenter Server system. Every associated host has a host proxy switch that contains the networking settings for the host that are configured on distributed switch.

For example, suppose you associate ESXi A and ESXi B hosts to a distributed switch and connect physical NIC vmnic1 of both hosts to uplink 1 on the switch. As a result, vmnic1 of hosts ESXi A and ESXi B is connected to uplink 1 on the distributed switch. On the host proxy switches of both hosts, physical NIC vmnic1 is connected to uplink port 1.

To ensure efficient use of host resources, the number of distributed ports of proxy switches are dynamically scaled up and down on hosts running ESXi 5.5 and later. A proxy switch on such a host can expand up to the maximum number of ports supported on the host. The port limit is determined based on the maximum number of virtual machines that the host can handle.

A distributed switch has one or more distributed port groups. You use distributed port groups to provide networking connectivity to virtual machines and to accommodate VMkernel traffic. You identify each distributed port group by using a network label, which must be unique to the current datacenter. A copy of every distributed port group that you create is also available on the host proxy switches of all hosts that are associated with the distributed switch. The policies that you configure to a distributed port group are consistent for all hosts in the distributed switch.

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

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, see the Cisco Systems Web site.

Create a vSphere Distributed Switch with the vSphere Web Client

Create a vSphere distributed switch on a datacenter to handle the networking configuration of multiple hosts at a time from a central place.

Procedure

1 In the vSphere Web Client, navigate to a datacenter.

2 In the navigator, right-click the datacenter and select **New Distributed Switch**.

3 In **Name and Location**, type a name for the new distributed switch, or accept the generated name, and click **Next.**

4 In **Select version**, select a distributed switch version and click **Next**.

Option	Description
Distributed Switch: 5.5.0	Compatible with ESXi 5.5 and later.
Distributed Switch: 5.1.0	Compatible with VMware ESXi 5.1 and later. Features released with later vSphere distributed switch versions are not supported.
Distributed Switch: 5.0.0	Compatible with VMware ESXi 5.0 and later. Features released with later vSphere distributed switch versions are not supported.
Distributed Switch: 4.1.0	Compatible with ESX/ESXi version 4.1 and later. Features released with later vSphere distributed switch versions are not supported.
Distributed Switch: 4.0.0	Compatible with ESX/ESXi version 4.0 and later. Features released with later vSphere distributed switch versions are not supported.

5 In **Edit Settings** configure the distributed switch settings.

   a Use the arrow buttons to select the **Number of uplinks**.

   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.

   b Use the drop-down menu to enable or disable **Network I/O Control**.

   By using Network I/O Control you can prioritize the access to network resources for certain types of infrastructure and workload traffic according to the requirements of your deployment. Network I/O Control continuously monitors the I/O load over the network and dynamically allocates available resources.

   c Select the **Create a default port group** check box to create a new distributed port group with default settings for this switch.

   d (Optional) To create a default distributed port group, type the port group name in the **Port group name**, or accept the generated name.

   If your system has custom port group requirements, create distributed port groups that meet those requirements after you add the distributed switch.

   e Click **Next**.

6 In **Ready to complete**, review the settings you selected and click **Finish**.

   Use the **Back** button to edit any settings.

A distributed switch is created on the datacenter. You can view the features supported on the distributed switch as well as other details by navigating to the new distributed switch and clicking the **Summary** tab.

What to do next

Add hosts to the distributed switch and configure their network adapters on the switch.

Upgrade a vSphere Distributed Switch to a Later Version with the vSphere Web Client

You can upgrade vSphere Distributed Switch version 4.0, 4.1, 5.0, or 5.1 to a later version. The upgrade lets the distributed switch take advantage of features that are available only in the later version.

The upgrade of a distributed switch is a non-disruptive operation, that is, the hosts and virtual machines attached to the switch do not experience any downtime.

NOTE To be able to restore the connectivity of the virtual machines and VMkernel adapters if the upgrade fails, back up the configuration of the distributed switch.

You can export the switch configuration before you upgrade vCenter Server if you upgrade from vCenter Server 5.1. If you upgrade vCenter Server from a version earlier than 5.1, back up the switch configuration after you upgrade vCenter Server to version 5.5.

If the upgrade is not successful, to recreate the switch with its port groups and connected hosts, you can import the switch configuration file with the **Preserve original distributed switch and port group identifiers** option selected in the Import Distributed Switch wizard.

See "Export vSphere Distributed Switch Configurations with the vSphere Web Client," on page 52 and "Import a vSphere Distributed Switch Configuration by Using the vSphere Web Client," on page 53.

Prerequisites

- Upgrade vCenter Server to the version 5.5.

- Upgrade all hosts connected to the distributed switch to ESXi 5.5.

Procedure

1 In the vSphere Web Client, navigate to the distributed switch.

2 Right-click the distributed switch and select **Upgrade Distributed Switch**.

3 Select the vSphere Distributed Switch version that you want to upgrade the switch to and click **Next**.

Option	Description
Version 5.5.0	Compatible with ESXi version 5.5 and later.
Version 5.1.0	Compatible with ESXi version 5.1 and later. Features released with later vSphere Distributed Switch versions are not supported.
Version 5.0.0	Compatible with ESXi version 5.0 and later. Features released with later vSphere Distributed Switch versions are not supported.
Version 4.1.0	Compatible with ESX/ESXi version 4.1 and later. Features released with later vSphere Distributed Switch versions are not supported.

4 Review host compatibility and click **Next**.

Some VMware ESX instances that are running on the distributed switch might be incompatible with the selected upgrade version. Upgrade or remove incompatible hosts, or select another upgrade version for the distributed switch.

5 Review your settings and click **Finish.**

After you upgrade the vSphere Distributed Switch, you cannot revert it to an earlier version. You cannot add VMware ESX hosts that are running an earlier incompatible version with the new switch version.

Edit General and Advanced vSphere Distributed Switch Settings in the vSphere Web Client

General settings for a vSphere Distributed Switch include the switch name and number of uplinks. Advanced settings for a distributed switch include Cisco Discovery Protocol and the maximum MTU for the switch.

Procedure

1 In the vSphere Web Client, navigate to the distributed switch.

2 Click **Manage** tab, click **Settings**, and select **Properties**.

3 Click **Edit**.

4 Click **General** to edit the vSphere Distributed Switch settings.

Option	Description
Name	Type the name for the distributed switch.
Number of uplinks	Select the number of uplink ports for the distributed switch. Click **Edit Uplink Names** to change the names of the uplinks.
Number of ports	The number of ports for this distributed switch. This cannot be edited.
Network I/O Control	Use the drop-down menu to enable or disable Network I/O control.
Description	Add or modify a description of the distributed switch settings.

5 Click **Advanced** to edit the vSphere Distributed Switch settings.

Option	Description
MTU (Bytes)	Maximum MTU size for the vSphere Distributed Switch. To enable jumbo frames, set a value greater than 1500 bytes.
Discovery Protocol	a Select Cisco Discovery Protocol, Link Layer Discovery Protocol, or disabled from the **Type** drop-down menu. b Set Operation to Listen, Advertise, or Both. For information about Discovery Protocol, see "Switch Discovery Protocol," on page 159.
Administrator Contact	Type the name and other details of the administrator for the distributed switch.

6 Click **OK**.

Managing Networking on Multiple Hosts on a vSphere Distributed Switch

You create and manage virtual networks on a vSphere Distributed Switch by adding hosts to the switch and connecting their network adapters to the switch. To create uniform networking configuration throughout multiple hosts on the distributed switch, you can use a host as a template and apply its configuration to other hosts.

- Tasks for Managing Host Networking on a vSphere Distributed Switch on page 29

  You can add new hosts to a vSphere Distributed Switch, connect network adapters to the switch, and remove hosts from the switch. In a production environment, you might need to keep the network connectivity up for virtual machines and VMkernel services while you manage host networking on the distributed switch.

- Add Hosts to a vSphere Distributed Switch in the vSphere Web Client on page 30

  To manage the networking of your vSphere environment by using a vSphere Distributed Switch, you must associate hosts with the switch. You connect the physical NICs, VMkernel adapters, and virtual machine network adapters of the hosts to the distributed switch.

- Configure Physical Network Adapters on a vSphere Distributed Switch in the vSphere Web Client on page 31

  For hosts that are associated with a distributed switch, you can assign physical NICs to uplinks on the switch. You can configure physical NICs on the distributed switch for multiple hosts at a time.

- Migrate VMkernel Adapters to a vSphere Distributed Switch in the vSphere Web Client on page 32

  Migrate VMkernel adapters to a distributed switch if you want to handle the traffic for VMkernel services by using only this switch and you no longer need the adapters on other standard or distributed switches.

- Create a VMkernel Adapter on a vSphere Distributed Switch in the vSphere Web Client on page 33

  Create a VMkernel adapter on hosts associated with a distributed switch to provide network connectivity to the hosts and to handle the traffic for vSphere vMotion, IP storage, Fault Tolerance logging, and Virtual SAN. You can create VMkernel adapters on multiple hosts simultaneously by using the Add and Manage Hosts wizard.

- Migrate Virtual Machine Networking to the vSphere Distributed Switch in the vSphere Web Client on page 34

  To manage virtual machine networking by using a distributed switch, migrate virtual machine network adapters to labeled networks on the switch.

- Use a Host as a Template to Create a Uniform Networking Configuration on a vSphere Distributed Switch in the vSphere Web Client on page 35

  If you plan to have hosts with a uniform networking configuration, you can select a host as a template and apply its configuration for physical NICs and VMkernel adapters to other hosts on the distributed switch.

- Remove Hosts from a vSphere Distributed Switch in the vSphere Web Client on page 36

  Remove hosts from a vSphere distributed switch if you have configured a different switch for the hosts.

Tasks for Managing Host Networking on a vSphere Distributed Switch

You can add new hosts to a vSphere Distributed Switch, connect network adapters to the switch, and remove hosts from the switch. In a production environment, you might need to keep the network connectivity up for virtual machines and VMkernel services while you manage host networking on the distributed switch.

Adding Hosts to a vSphere Distributed Switch

Consider preparing your environment before you add new hosts to a distributed switch.

- Create distributed port groups for virtual machine networking.
- Create distributed port groups for VMkernel services. For example, create distributed port groups for management network, vMotion, and Fault Tolerance.
- Configure enough uplinks on the distributed switch for all physical NICs that you want to connect to the switch. For example, if the hosts that you want to connect to the distributed switch have eight physical NICs each, configure eight uplinks on the distributed switch.
- Make sure that the configuration of the distributed switch is prepared for services with specific networking requirements. For example, iSCSI has specific requirements for the teaming and failover configuration of the distributed port group where you connect the iSCSI VMkernel adapter.

You can use the Add and Manage Hosts wizard in the vSphere Web Client to add multiple hosts at a time.

Managing Network Adapters on a vSphere Distributed Switch

After you add hosts to a distributed switch, you can connect physical NICs to uplinks on the switch, configure virtual machine network adapters, and manage VMkernel networking.

If some hosts on a distributed switch are associated to other switches in your datacenter, you can migrate network adapters to or from the distributed switch.

If you migrate virtual machine network adapters or VMkernel adapters, make sure that the destination distributed port groups have at least one active uplink, and the uplink is connected to a physical NIC on the hosts. Another approach is to migrate physical NICs, virtual network adapters, and VMkernel adapters simultaneously.

If you migrate physical NICs, leave at least one active NIC that handles the traffic of port groups. For example, if *vmnic0* and *vmnic1* handle the traffic of the *VM Network* port group, migrate *vmnic0* and leave *vmnic1* connected to the group.

Removing Hosts from a vSphere Distributed Switch

Before you remove hosts from a distributed switch, you must migrate the network adapters that are in use to a different switch.

- To add hosts to a different distributed switch, you can use the Add and Manage Hosts wizard to migrate the network adapters on the hosts to the new switch all together. You can then remove the hosts safely from their current distributed switch.

- To migrate host networking to standard switches, you must migrate the network adapters on stages. For example, remove physical NICs on the hosts from the distributed switch by leaving one physical NIC on every host connected to the switch to keep the network connectivity up. Next, attach the physical NICs to the standard switches and migrate VMkernel adapters and virtual machine network adapters to the switches. Lastly, migrate the physical NIC that you left connected to the distributed switch to the standard switches.

Add Hosts to a vSphere Distributed Switch in the vSphere Web Client

To manage the networking of your vSphere environment by using a vSphere Distributed Switch, you must associate hosts with the switch. You connect the physical NICs, VMkernel adapters, and virtual machine network adapters of the hosts to the distributed switch.

Prerequisites

- Verify that enough uplinks are available on the distributed switch to assign to the physical NICs that you want to connect to the switch.
- Verify that there is at least one distributed port group on the distributed switch.
- Verify that the distributed port group have active uplinks configured in its teaming and failover policy.

If you migrate or create VMkernel adapters for iSCSI, verify that the teaming and failover policy of the target distributed port group meets the requirements for iSCSI:

- Verify that only one uplink is active, the standby list is empty, and the rest of the uplinks are unused.
- Verify that only one physical NIC per host is assigned to the active uplink.

Procedure

1 In the vSphere Web Client, navigate to the distributed switch.

2 From the **Actions** menu, select **Add and Manage Hosts**.

3 Select **Add hosts** and click **Next**.

4 Click **New hosts**, select from the hosts in your datacenter, and click **OK**.

5 Select the tasks for configuring network adapters to the distributed switch and click **Next**.

6 Configure physical NICs on the distributed switch.

   a From the On other switches/unclaimed list, select a physical NIC.

   If you select physical NICs that are already connected to other switches, they are migrated to the current distributed switch.

   b Click **Assign uplink**.

   c Select an uplink and click **OK**.

   For consistent network configuration, you can connect one and the same physical NIC on every host to the same uplink on the distributed switch.

   For example, if you are adding two hosts connect *vmnic1* on of each host to *Uplink1* on the distributed switch.

7 Click **Next**.

8 Configure VMkernel adapters.

   a Select a VMkernel adapter and click **Assign port group**.

   b Select a distributed port group and click **OK**.

9 Review the impacted services as well as the level of impact.

Option	Description
No impact	The service will continue its normal function after the new networking configuration is applied.
Important impact	The normal function of the service might be disrupted if the new networking configuration is applied.
Critical impact	The normal function of the service will be interrupted if the new networking configuration is applied.

   a If impact on a service is important or critical, click the service and review the reasons that are displayed in the Analysis details pane.

   b After you troubleshoot the impact on all dependent services, proceed with your networking configuration.

10 Click **Next**.

11 Configure virtual machine networking.

   a To connect all network adapters of a virtual machine to a distributed port group, select the virtual machine, or select an individual network adapter to connect only that adapter.

   b Click **Assign port group**.

   c Select a distributed port group from the list and click **OK**.

12 Click **Next** and click **Finish**.

What to do next

Having hosts associated with the distributed switch, you can manage physical NICs, VMkernel adapters, and virtual machine network adapters.

Configure Physical Network Adapters on a vSphere Distributed Switch in the vSphere Web Client

For hosts that are associated with a distributed switch, you can assign physical NICs to uplinks on the switch. You can configure physical NICs on the distributed switch for multiple hosts at a time.

For consistent networking configuration throughout all hosts, you can assign the same physical NIC on every host to the same uplink on the distributed switch. For example, you can assign *vmnic1* from hosts *ESXi A* and *ESXi B* to *Uplink 1*.

Procedure

1 In the vSphere Web Client, navigate to the distributed switch.

2 From the **Actions** menu, select **Add and Manage Hosts**.

3 Select **Manage host networking** and click **Next**.

4 Click **Attached hosts** and select from the hosts that are associated with the distributed switch.

5 Click **Next**.

6 Select **Manage physical adapters** and click **Next**.

7 From the On other switches/unclaimed list select a physical NIC .

   If you select physical NICs that are already assigned to other switches, they are migrated to the current distributed switch.

8 Click **Assign uplink**.

9 Select an uplink or select **Auto-assign**.

10 Click **Next**.

11 Review the impacted services as well as the level of impact.

Option	Description
No impact	The service will continue its normal function after the new networking configuration is applied.
Important impact	The normal function of the service might be disrupted if the new networking configuration is applied.
Critical impact	The normal function of the service will be interrupted if the new networking configuration is applied.

a If impact on a service is important or critical, click the service and review the reasons that are displayed in the Analysis details pane.

b After you troubleshoot the impact on all dependent services, proceed with your networking configuration.

12 Click **Next** and click **Finish**.

Migrate VMkernel Adapters to a vSphere Distributed Switch in the vSphere Web Client

Migrate VMkernel adapters to a distributed switch if you want to handle the traffic for VMkernel services by using only this switch and you no longer need the adapters on other standard or distributed switches.

Procedure

1 In the vSphere Web Client, navigate to the distributed switch.

2 From the **Actions** menu, select **Add and Manage Hosts**.

3 Select **Manage host networking** and click **Next**.

4 Click **Attached hosts** and select from the hosts that are associated with the distributed switch.

5 Click **Next**.

6 Select **Manage VMkernel adapters** and click **Next**.

7 Select the adapter and click **Assign port group**.

8 Select a distributed port group and click **OK**.

9 Click **Next**.

10 Review the impacted services as well as the level of impact.

Option	Description
No impact	The service will continue its normal function after the new networking configuration is applied.
Important impact	The normal function of the service might be disrupted if the new networking configuration is applied.
Critical impact	The normal function of the service will be interrupted if the new networking configuration is applied.

a If impact on a service is important or critical, click the service and review the reasons that are displayed in the Analysis details pane.

b After you troubleshoot the impact on all dependent services, proceed with your networking configuration.

11 Click **Next** and click **Finish**.

Create a VMkernel Adapter on a vSphere Distributed Switch in the vSphere Web Client

Create a VMkernel adapter on hosts associated with a distributed switch to provide network connectivity to the hosts and to handle the traffic for vSphere vMotion, IP storage, Fault Tolerance logging, and Virtual SAN. You can create VMkernel adapters on multiple hosts simultaneously by using the Add and Manage Hosts wizard.

You should dedicate one distributed port group for each VMkernel adapter. One VMkernel adapter should handle only one traffic type.

Procedure

1 In the vSphere Web Client, navigate to the distributed switch.

2 From the **Actions** menu, select **Add and Manage Hosts**.

3 Select **Manage host networking** and click **Next**.

4 Click **Attached hosts** and select from the hosts that are associated with the distributed switch.

5 Click **Next**.

6 Select **Manage VMkernel adapters** and click **Next**.

7 Click **New adapter**.

The Add Networking wizard opens.

8 On the Select target device page of the Add Networking wizard, select a distributed port group.

9 On the Port properties page, configure the settings for the VMkernel adapter.

Option	Description
Network label	The network label is inherited from the label of the distributed port group.
IP settings	Select IPv4, IPv6, or both. NOTE The IPv6 option does not appear on hosts that do not have IPv6 enabled.
TCP/IP stack	If custom stacks are available, select one from the list.
Enable services	You can enable services for the default TCP/IP stack on the host. Select from the available services: ■ **vMotion traffic**. Enables the VMkernel adapter to advertise itself to another host as the network connection where vMotion traffic is sent. You can enable this property for only one vMotion and IP storage VMkernel adapter per host. If this property is not enabled for any VMkernel adapter, migration with vMotion to the selected host is not possible. ■ **Fault Tolerance traffic**. Enables Fault Tolerance logging on the host. ■ **Management traffic**. Enables the management traffic for the host and vCenter Server. Typically, hosts have such a VMkernel adapter created when the ESXi software is installed. You can create another VMkernel adapter for management traffic on the host to provide redundancy. ■ **Virtual SAN**. Enables the Virtual SAN traffic on the host. Every host that is part of a Virtual SAN cluster must have such a VMkernel adapter.

10 (Optional) On the IPv4 settings page, select an option for obtaining IP addresses.

Option	Description
Obtain IP settings automatically	Use DHCP to obtain IP settings.
Use static IP settings	Type the IPv4 IP address and subnet mask for the VMkernel adapter. The VMkernel Default Gateway and DNS server addresses for IPv4 are obtained from the selected TCP/IP stack.

11 (Optional) On the IPv6 settings page, select an option for obtaining IPv6 addresses.

Option	Description
Obtain IPv6 addresses automatically through DHCP	Use DHCP to obtain IPv6 addresses.
Obtain IPv6 addresses automatically through Router Advertisement	Use router advertisement to obtain IPv6 addresses.
Static IPv6 addresses	a Click **Add** to add a new IPv6 address. b Type the IPv6 address and subnet prefix length, and click **OK**. c To change the VMkernel default gateway, click **Edit**.

12 Review your setting selections in the Ready to complete page and click **Finish**.

13 Follow the prompts to complete the wizard.

Migrate Virtual Machine Networking to the vSphere Distributed Switch in the vSphere Web Client

To manage virtual machine networking by using a distributed switch, migrate virtual machine network adapters to labeled networks on the switch.

Prerequisites

Verify that at least one distributed port group intended for virtual machine networking exists on the distributed switch.

Procedure

1 In the vSphere Web Client, navigate to the distributed switch.

2 From the **Actions** menu, select **Add and Manage Hosts**.

3 Select **Manage host networking** and click **Next**.

4 Click **Attached hosts** and select from the hosts that are associated with the distributed switch.

5 Click **Next**.

6 Select **Migrate virtual machine networking** and click **Next**.

7 Configure virtual machine network adapters to the distributed switch.

   a To connect all network adapters of a virtual machine to a distributed port group, select the virtual machine, or select an individual network adapter to connect only that adapter.

   b Click **Assign port group**.

   c Select a distributed port group from the list and click **OK**.

8 Click **Next** and click **Finish**.

Use a Host as a Template to Create a Uniform Networking Configuration on a vSphere Distributed Switch in the vSphere Web Client

If you plan to have hosts with a uniform networking configuration, you can select a host as a template and apply its configuration for physical NICs and VMkernel adapters to other hosts on the distributed switch.

Procedure

1 In the vSphere Web Client, navigate to the distributed switch.

2 From the **Actions** menu, select **Add and Manage Hosts**.

3 Select a task for managing host networking and click **Next**.

4 Select the hosts to add or manage on the distributed switch.

5 At the bottom of the dialog box, select **Configure identical networking settings on multiple hosts** and click **Next**.

6 Select a host to use as a template and click **Next**.

7 Select the network adapter tasks and click **Next**.

8 On the Manage physical network adapters and Manage VMkernel network adapters pages, make the configuration changes that you need on the template host, and click **Apply to all** for all other hosts.

9 On the Ready to complete page, click **Finish**.

Example: Configure Physical NICs by Using a Template Host

Connect physical NICs on two hosts simultaneously to a vSphere distributed switch by using the template host mode in the Add and Manage Hosts wizard. On the Manage physical network adapters page of the wizard, assign two physical NICs to uplinks on the template host, and then click **Apply to all** to create the same configuration on the other host.

Figure 4-2. Applying Physical NICs Configuration on a vSphere Distributed Switch by Using a Template Host

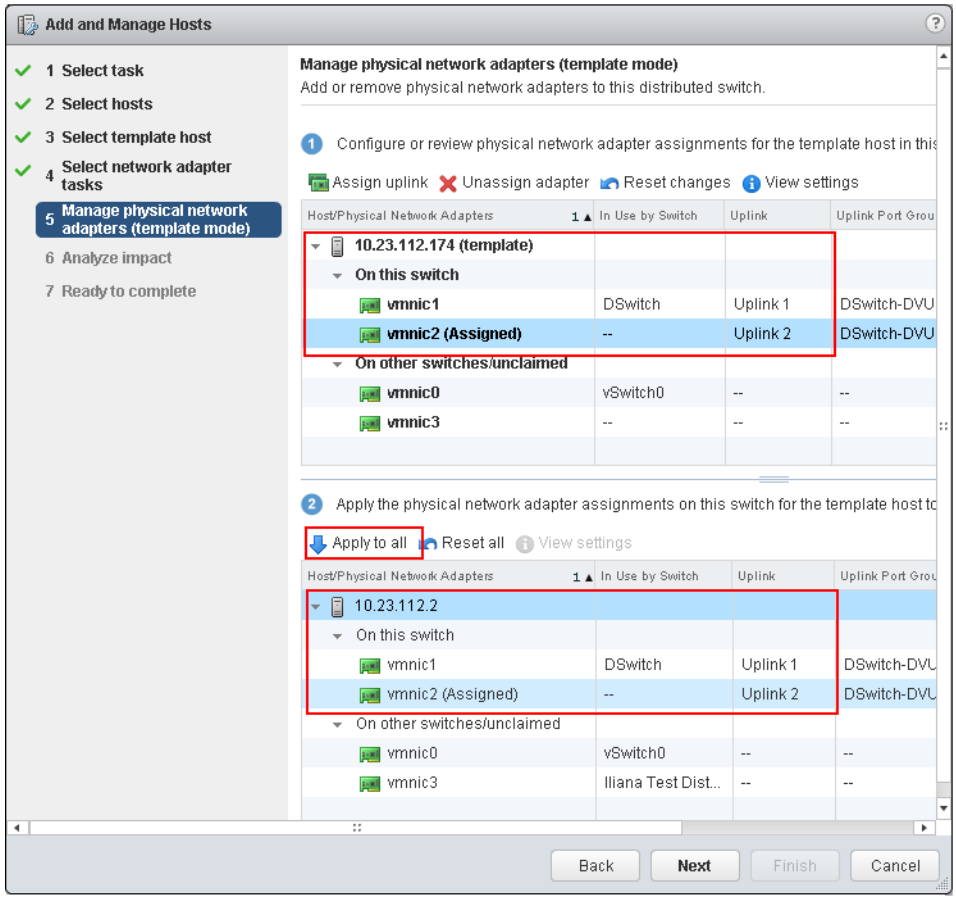

Remove Hosts from a vSphere Distributed Switch in the vSphere Web Client

Remove hosts from a vSphere distributed switch if you have configured a different switch for the hosts.

Prerequisites

- Verify that physical NICs on the target hosts are migrated to a different switch.
- Verify that VMkernel adapters on the hosts are migrated to a different switch.
- Verify that virtual machine network adapters are migrated to a different switch.

For details about migrating network adapters to different switches, see "Tasks for Managing Host Networking on a vSphere Distributed Switch," on page 29

Procedure

1 In the vSphere Web Client, navigate to the distributed switch.

2 From the **Actions** menu, select **Add and Manage Hosts**.

3 Select **Remove hosts** and click **Next**.

4 Select the hosts you want to remove and click **Next**.

5 Click **Finish**.

Managing Networking on Host Proxy Switches

You can change the configuration of the proxy switch on every host that is associated with a vSphere distributed switch. You can manage physical NICs, VMkernel adapters, and virtual machine network adapters.

For details about setting up VMkernel networking on host proxy switches, see "Create a VMkernel Adapter on a vSphere Distributed Switch in the vSphere Web Client," on page 33.

Migrate Network Adapters on a Host to a vSphere Distributed Switch in the vSphere Web Client

For hosts associated with a distributed switch, you can migrate network adapters from a standard switch to the distributed switch. You can migrate physical NICs, VMkernel adapters, and virtual machine network adapters at the same time.

If you want to migrate virtual machine network adapters or VMkernel adapters, make sure that the destination distributed port groups have at least one active uplink, and the uplink is connected to a physical NIC on this host. Alternatively, migrate physical NICs, virtual network adapters, and VMkernel adapters at once.

If you want to migrate physical NICs, make sure that the source port groups on the standard switch have at least one physical NIC to handle their traffic. For example, if you migrate a physical NIC that is assigned to a port group for virtual machine networking, make sure that the port group is connected to at least one physical NIC. Otherwise the virtual machines on same VLAN on the standard switch will have connectivity between each other but not to the external network.

Procedure

1 In the vSphere Web Client, navigate to the host.

2 On the **Manage** tab, click **Networking**, and select **Virtual switches**.

3 Select the destination distributed switch and click **Migrate physical or virtual network adapters**.

4 Select the tasks for migrating network adapters and click **Next**.

5 Configure physical NICs.

   a From the **On other switches/unclaimed** list, select a physical NIC and click **Assign uplink**.

   b Select an uplink and click **OK**.

   c Click **Next**.

6 Configure VMkernel adapters.

   a Select an adapter and click **Assign port group**.

   b Select a distributed port group and click **OK**.

   You should connect one VMkernel adapter to one distributed port group at a time.

   c Click **Next**.

7 Review the services that are affected from the new networking configuration.

   a If there is an important or serious impact reported on a service, click the service and review the analysis details.

      For example, an important impact on iSCSI might be reported as a result from an incorrect teaming and failover configuration on the distributed port group where you migrate the iSCSI VMkernel adapter. You must leave one active uplink on the teaming and failover order of the distributed port group, leave the standby list empty, and move the rest of the uplinks to unused.

   b After troubleshooting any impact on the affected services, click **Next**.

8 Configure virtual machine network adapters.

   a Select a virtual machine or a virtual machine network adapter and click **Assign port group**.

      If you select a virtual machine, you migrate all network adapters on the virtual machine. If you select a network adapter, you migrate only this network adapter.

   b Select a distributed port group from the list anc click **OK**.

   c Click **Next**.

9 On the Ready to complete page review the new networking configuration and click **Finish**.

Migrate a VMkernel Adapter on a Host to a vSphere Standard Switch in the vSphere Web Client

If a host is associated with a distributed switch, you can migrate VMkernel adapters from the distributed to a standard switch.

For details about creating VMkernel adapters on a vSphere distributed switch, see "Create a VMkernel Adapter on a vSphere Distributed Switch in the vSphere Web Client," on page 33.

Procedure

1 In the vSphere Web Client, navigate to the host.

2 On the **Manage** tab, click **Networking**, and select **Virtual switches**.

3 Select the destination standard switch from the list.

4 Click **Migrate a VMkernel adapter**.

5 On the Select VMkernel network adapter page, select the virtual network adapter to migrate to the standard switch from the list.

6 On the Configure settings page, edit the **Network Label** and **VLAN ID** for the network adapter.

7 Review the migration details on the Ready to complete page and click **Finish**.

   Click **Back** to edit settings.

Assign a Physical NIC of a Host to a vSphere Distributed Switch in the vSphere Web Client

You can assign physical NICs of a host that is associated with a distributed switch to uplink port on the host proxy switch.

Procedure

1 In the vSphere Web Client, navigate to the host.

2 On the **Manage** tab, click **Networking**, and select **Virtual switches**.

3 Select a distributed switch from the list.

4 Click **Manage the physical network adapters**.

5 Select a free uplink from the list and click **Add adapter**.

6 Select a physical NIC and click **OK**.

Remove a Physical NIC from a vSphere Distributed Switch in the vSphere Web Client

You can remove a physical NIC of a host from an uplink on a vSphere distributed switch.

Procedure

1 In the vSphere Web Client, navigate to the host.

2 On the **Manage** tab, click **Networking**, and select **Virtual switches**.

3 Select the distributed switch.

4 Click **Manage the physical network adapters**.

5 Select an uplink an click **Remove selected adapters**.

6 Click **OK**.

What to do next

When you remove physical NICs from active virtual machines, you might see the NICs you removed reported in the vSphere Web Client. See "Removing NICs from Active Virtual Machines," on page 39.

Removing NICs from Active Virtual Machines

When you remove NICs from active virtual machines, you might still see the NICs you have removed in the vSphere Web Client.

Removing NICs from an Active Virtual Machine Without a Guest Operating System Installed

You cannot remove NICs from an active virtual machine on which no operating system is installed.

The vSphere Web Client might report that the NIC has been removed, but you continue to see it attached to the virtual machine.

Removing 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 vSphere Web Client for some time. If you click **Edit Settings** for the virtual machine, you might see the removed NIC listed even after the task is complete. The Edit Settings dialog box for the virtual machine does not immediately display the removed NIC.

You might also still see the NIC attached to the virtual machine if the guest operating system of the virtual machine does not support hot removal of NICs.

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 in the vSphere Web Client

Add a distributed port group to a vSphere distributed switch to create a distributed switch network for your virtual machines and to associate VMkernel adapters.

Procedure

1 Browse to a distributed switch in the vSphere Web Client.

2 Right-click the distributed switch in the navigator and select **New distributed port group**.

3 In the **Select name and location** section, type the name of the new distributed port group, or accept the generated name, and click **Next**.

4 In the **Configure settings** section, set the general properties for the new distributed port group and click **Next**.

Setting	Description
Port binding	Choose when ports are assigned to virtual machines connected to this distributed port group. ■ **Static binding**: Assign a port to a virtual machine when the virtual machine connects to the distributed port group. ■ **Dynamic binding**: 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 has been deprecated since ESXi 5.0. ■ **Ephemeral**: No port binding. You can assign a virtual machine to a distributed port group with ephemeral port binding also when connected to the host.
Port allocation	■ **Elastic:** The default number of ports is eight. When all ports are assigned, a new set of eight ports is created. This is the default. ■ **Fixed**: The default number of ports is set to eight. No additional ports are created when all ports are assigned.
Number of ports	Enter the number of ports on the distributed port group.
Network resource pool	Use the drop-down menu to assign the new distributed port group to a user-defined network resource pool. If you have not created a network resource pool, this menu is empty.
VLAN	Use the **Type** drop-down menu to select VLAN options: ■ **None**: Do not use VLAN. ■ **VLAN**: In the **VLAN ID** field, enter a number between 1 and 4094. ■ **VLAN Trunking**: Enter a VLAN trunk range. ■ **Private VLAN**: Select a private VLAN entry. If you did not create any private VLANs, this menu is empty.
Advanced	Select this check box to customize the policy configurations for the new distributed port group.

5 (Optional) In the **Security** section, edit the security exceptions and click **Next**.

Setting	Description
Promiscuous mode	■ **Reject**. Placing an adapter in promiscuous mode from the guest operating system does not result in receiving frames for other virtual machines. ■ **Accept**. If an adapter is placed in promiscuous mode from the guest operating system, the switch allows the guest adapter to receive all frames passed on the switch in compliance with the active VLAN policy for the port where the adapter is connected. Firewalls, port scanners, intrusion detection systems and so on, need to run in promiscuous mode.
MAC address changes	■ **Reject**. If you set this option to **Reject** and the guest operating system changes the MAC address of the adapter to a value different from the address in the `.vmx` configuration file, the switch drops all inbound frames to the virtual machine adapter. . If the guest operating system changes the MAC address back, the virtual machine receives frames again. ■ **Accept**. If the guest operating system changes the MAC address of a network adapter, the adapter receives frames to its new address.
Forged transmits	■ **Reject**. The switch drops any outbound frame with a source MAC address that is different from the one in the `.vmx` configuration file. ■ **Accept**. The switch does not perform filtering and permits all outbound frames.

6 (Optional) In the **Traffic shaping** section, enable or disable Ingress or Egress traffic shaping and click **Next**.

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

7 (Optional) In the **Teaming and failover** section, edit the settings and click **Next**.

Setting	Description
Load balancing	Specify how to choose an uplink. ■ **Route based on the originating virtual port**. Choose an uplink based on the virtual port where the traffic entered the distributed switch. ■ **Route based on IP hash**. Choose an uplink based on a hash of the source and destination IP addresses of each packet. For non-IP packets, whatever is at those offsets is used to compute the hash. ■ **Route based on source MAC hash**. Choose an uplink based on a hash of the source Ethernet. ■ **Route based on physical NIC load**. Choose an uplink based on the current loads of physical NICs. ■ **Use explicit failover order**. Always use the highest order uplink from the list of Active adapters which passes failover detection criteria. NOTE IP-based teaming requires that the physical switch be configured with etherchannel. For all other options, disable etherchannel.
Network failover detection	Specify the method to use for failover detection. ■ **Link Status only**. Relies solely on the link status that the network adapter provides. This option detects failures, such as cable pulls and physical switch power failures, but not configuration errors, such as a physical switch port being blocked by spanning tree or that is misconfigured to the wrong VLAN or cable pulls on the other side of a physical switch. ■ **Beacon Probing**. Sends out and listens for beacon probes on all NICs in the team and uses this information, in addition to link status, to determine link failure. This detects many of the failures previously mentioned that are not detected by link status alone. NOTE Do not use beacon probing with IP-hash load balancing.
Notify switches	Select **Yes** or **No** to notify switches in the case of failover. If you select **Yes**, whenever a virtual NIC is connected to the 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. NOTE Do not use this option when the virtual machines using the port group are using Microsoft Network Load Balancing in unicast mode. No such issue exists with NLB running in multicast mode.
Failback	Select **Yes** or **No** to disable or enable failback. This option determines how a physical adapter is returned to active duty after recovering from a failure. If failback is set to **Yes** (default), the adapter is returned to active duty immediately upon recovery, displacing the standby adapter that took over its slot, if any. If failback is set to **No**, a failed adapter is left inactive even after recovery until another currently active adapter fails, requiring its replacement.
Failover order	Specify how to distribute the work load for uplinks. To use some uplinks but reserve others for emergencies if the uplinks in use fail, set this condition by moving them into different groups: ■ **Active Uplinks**. Continue to use the uplink when the network adapter connectivity is up and active. ■ **Standby Uplinks** . Use this uplink if one of the active adapter's connectivity is down. ■ **Unused Uplinks** . Do not use this uplink. NOTE When using IP-hash load balancing, do not configure standby uplinks.

8 (Optional) In the **Monitoring** section, enable or disable NetFlow and click **Next**.

Setting	Description
Disabled	NetFlow is disabled on the distributed port group.
Enabled	NetFlow is enabled on the distributed port group. NetFlow settings can be configured at the vSphere distributed switch level.

9 (Optional) In the **Miscellaneous** section, select **Yes** or **No** and click **Next**.

Selecting **Yes** shuts down all ports in the port group. This action might disrupt the normal network operations of the hosts or virtual machines using the ports.

10 (Optional) In the **Edit additional settings** section, add a description of the port group and set any policy overrides per port and click **Next**.

11 Review your settings in the **Ready to complete** section and click **Finish**.

Click the **Back** button to change any settings.

Edit General Distributed Port Group Settings with the vSphere Web Client

You can edit general distributed port group settings such as the distributed port group name, port settings and network resource pool.

Procedure

1 Locate a distributed port group in the vSphere Web Client.

a Select a distributed switch and click the **Related Objects** tab.

b Click **Distributed Port Groups**.

2 Right-click the distributed port group and select **Edit settings**.

3 Select **General** to edit the following distributed port group settings.

Option	Description
Name	The name of distributed port group. You can edit the name in the text field.
Port binding	Choose when ports are assigned to virtual machines connected to this distributed port group. ■ **Static binding**: Assign a port to a virtual machine when the virtual machine connects to the distributed port group. ■ **Dynamic binding**: 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 has been deprecated since ESXi 5.0. ■ **Ephemeral**: No port binding. You can also assign a virtual machine to a distributed port group with ephemeral port binding when connected to the host.
Port allocation	■ **Elastic:** The default number of ports is set to eight. When all ports are assigned, a new set of eight ports is created. This is the default. ■ **Fixed**: The default number of ports is set to eight. No additional ports are created when all ports are assigned.
Number of ports	Enter the number of ports on the distributed port group.
Network resource pool	Use the drop-down menu to assign the new distributed port group to a user-defined network resource pool. If you have not created a network resource pool, this menu is empty.
Description	Enter any information about the distributed port group in the description field.

4 Click **OK**.

Edit Advanced Distributed Port Group Settings with the vSphere Web Client

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

Procedure

1 Locate a distributed port group in the vSphere Web Client.

   a To locate a distributed port group, select a distributed switch and click the **Related Objects** tab.

   b Click **Distributed Port Groups** and select a distributed port group from the list.

2 Click the **Manage** tab and click **Settings**.

3 Click **Edit**.

4 Select the**Advanced** page to edit the distributed port group settings.

Option	Description
Configure reset at disconnect	From the drop-down menu, enable or disable reset at disconnect. 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.
Override port policies	Select the distributed port group policies to be overridden on a per-port level.

5 (Optional) Use the policy pages to set overrides for each port policy.

6 Click **OK**.

Remove a Distributed Port Group in the vSphere Web Client

Remove a distributed port group when you no longer need the corresponding labeled network to provide connectivity for virtual machines or VMkernel networking.

Prerequisites

- Verify that all virtual machines connected to the corresponding labeled network are migrated to a different labeled network.
- Verify that all VMkernel adapters connected to the distributed port group are migrated to a different port group, or are deleted.

Procedure

1 In the vSphere Web Client, navigate to the distributed port group.

2 From the **Actions** menu, select **All vCenter Actions > Remove from Inventory**.

Export, Import, and Restore vSphere Distributed Port Group Configurations

You can export vSphere distributed port group configurations to a file. The configuration file allows you to preserve valid port group configurations, enabling distribution of these configurations to other deployments.

You can export port group information at the same time you export distributed switch configurations. See "Export, Import, and Restore Distributed Switch Configurations," on page 52.

Export vSphere Distributed Port Group Configurations with the vSphere Web Client

You can export vSphere distributed port group configurations to a file. The configuration preserves valid network configurations, enabling distribution of these configurations to other deployments.

This functionality is available only with the vSphere Web Client 5.1 or later. However, you can export settings from any version of a distributed port if you use the vSphere Web Client 5.1 or later.

Procedure

1 Locate a distributed port group in the vSphere Web Client.

   a To locate a distributed port group, select a distributed switch and click the **Related Objects** tab.

   b Click **Distributed Port Groups** and select a distributed port group from the list.

2 Right-click the distributed port group in the navigator and select **All vCenter Actions > Export Configuration**.

3 (Optional) Type notes about this configuration in the **Descriptions** field.

4 Click **OK**.

   Click **Yes** to save the configuration file to your local system.

You now have a configuration file that contains all the settings for the selected distributed port group. You can use this file to create multiple copies of this configuration on an existing deployment, or overwrite settings of existing distributed port groups to conform to the selected settings.

What to do next

You can use the exported configuration file to do the following tasks:

- To create a copy of the exported distributed port group, see "Import a vSphere Distributed Port Group Configuration with the vSphere Web Client," on page 45.
- To overwrite settings on an existing distributed port group, see "Restore a vSphere Distributed Port Group Configuration with the vSphere Web Client," on page 46.

Import a vSphere Distributed Port Group Configuration with the vSphere Web Client

Use import to create a distributed port group from a configuration file. Any existing distributed port groups are converted to conform to the settings in the configuration file.

This functionality is available only with the vSphere Web Client 5.1 or later. However, you can export settings from any version of distributed port if you use the vSphere Web Client 5.1 or later.

Procedure

1 Browse to a distributed switch in the vSphere Web Client.

2 Right-click the distributed switch in the navigator and select **All vCenter Actions > Import Distributed Port Group**.

3 Browse to the location of your saved configuration file and click **Next**.

   You can use a distributed port group configuration file, or a distributed switch configuration file. However, you can use a file containing both distributed switch and distribute port group configurations only if the file contains settings for a single port group. If multiple port group settings are saved in the distributed switch configuration file, you must use a different file.

4 Review the import settings before completing the import.

5 Click **Finish.**

Restore a vSphere Distributed Port Group Configuration with the vSphere Web Client

Use the restore option to reset the configuration of an existing distributed port group to the settings in a configuration file.

This functionality is available only with the vSphere Web Client 5.1 or later. However, you can restore settings from any version of distributed switch if you use the vSphere Web Client 5.1 or later.

Procedure

1 Locate a distributed port group in the vSphere Web Client.

   a To locate a distributed port group, select a distributed switch and click **Related Objects**.

   b Click **Distributed Port Group** and select a distributed port group from the list.

2 Right-click the distributed port group in the navigator and select **All vCenter Actions > Restore Configuration**.

3 Select one of the following and click **Next**:

   - **Restore to a previous configuration** to roll your port group configuration back one step. You cannot restore the port group configuration completely if you have performed more than one step.
   - **Restore configuration from a file** lets you restore the port group configuration from an exported backup file. You can also use a distributed switch backup file as long as it contains configuration information for the port group.

4 Review the summary information for the restore.

   The restore operation overwrites the current settings of the distributed port group with the settings from the backup. If you are restoring the port group configuration from a switch backup file, the restore operation does not delete existing port groups that are not a part of the file.

5 Click **Finish**.

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 with the vSphere Web Client

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

Procedure

1 Locate a distributed port group in the vSphere Web Client.

   a To locate a distributed port group, select a distributed switch and click the **Related Objects** tab.

   b Click **Distributed Port Groups** and select a distributed port group from the list.

2 Click the **Manage** tab, and click **Ports**.

3 Click **Start Monitoring Port State**.

The ports table for the distributed port group displays runtime statistics for each distributed port.

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

Option	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 with the vSphere Web Client

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

Procedure

1 Locate a distributed port group in the vSphere Web Client.

a To locate a distributed port group, select a distributed switch and click the **Related Objects** tab.

b Click **Distributed Port Groups** and select a distributed port group from the list.

2 Click the **Manage** tab, and click **Ports**.

3 Select a distributed port from the table.

Information about the distributed port appears at the bottom of the screen.

4 Click **Edit distributed port settings**.

5 On the **Properties** page and policy pages, edit information about the distributed port and click **OK**.

If overrides are not allowed, the policy options are dimmed.

You can allow overrides at the port level by changing the **Advanced** settings of the distributed port group. See "Edit Advanced Distributed Port Group Settings with the vSphere Web Client," on page 44.

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 with the vSphere Web Client

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 Browse to a datacenter in the vSphere Web Client navigator.

2 Right-click the datacenter in the navigator and select **Migrate VM to Another Network**.

3 Select a source network.

- Select **Specific network** and use the **Browse** button to select a specific source network.
- Select **No network** to migrate all virtual machine network adapters that are not connected to any other network.

4 Select a destination network. Use **Browse** to select a specific destination network and click **Next**.

5 Select virtual machines from the list to migrate from the source network to the destination network and click **Next**.

6 Review your selections and click **Finish**.

Click **Back** to edit any selections.

Connect an Individual Virtual Machine to a Distributed Port Group in the vSphere Web Client

Connect an individual virtual machine to a vSphere distributed switch by modifying the NIC configuration of the virtual machine.

Procedure

1 Locate the virtual machine in the vSphere Web Client.

a Select a datacenter, folder, cluster, resource pool, or host and click the **Related Objects** tab.

b Click **Virtual Machines** and select the virtual machine from the list.

2 On the **Manage** tab of the virtual machine, select **Settings > VM Hardware**.

3 Click **Edit**.

4 Expand the **Network adapter** section and select a distributed port group from the drop-down menu.

5 Click **OK**.

Topology Diagrams of a vSphere Distributed Switch in the vSphere Web Client

The topology diagrams of a vSphere Distributed Switch in the vSphere Web Client show the structure of virtual machine adapters, VMkernel adapters, and physical adapters in the switch.

You can examine the components, arranged in port groups, whose traffic is handled by the switch, and the connections between them. The diagram displays information about the physical adapter that connects the virtual adapters to the external network.

You can view the components that are running on the entire virtual distributed switch and on each host participating in it.

Central Topology Diagram

You can use the central topology diagram of the switch to locate and edit the settings for distributed port groups and uplink groups associated with multiple hosts. You can initiate migration of virtual machine adapters from a port group to a destination on the same or different switch. You can also reorganize the hosts and their networking on the switch by using the Add and Manage Hosts wizard.

Topology Diagram of a Host Proxy Switch

The topology diagram of a host proxy switch shows the adapters attached to the switch ports on the host. You can edit the settings of the VMkernel and physical adapters.

Diagram Filters

You can use diagram filters to limit the information displayed in topology diagrams. The default filter limits the topology diagram to display 32 port groups, 32 hosts, and 1024 virtual machines.

You can change the scope of the diagram by using no filters or by applying custom filters. By using a custom filter, you can view information only about a set of virtual machines, a set of port groups on certain hosts, or a port. You can create filters from the central topology diagram of the distributed switch.

View the Topology of a vSphere Distributed Switch in the vSphere Web Client

Examine the organization of components that are connected to the distributed switch across the hosts in a vCenter Server.

Procedure

1 Navigate to the vSphere distributed switch in the vSphere Web Client.

2 On the **Manage** tab, click **Settings** and select **Topology**.

By default the diagram shows up to 32 distributed port groups, 32 hosts, and 1024 virtual machines.

Example: Diagram of a Distributed Switch That Connects the VMkernel and Virtual Machines to the Network

In your virtual environment, a vSphere Distributed Switch handles VMkernel adapters for vSphere vMotion and for the management network, and virtual machines grouped. You can use the central topology diagram to examine whether a virtual machine or VMkernel adapter is connected to the external network and to identify the physical adapter that carries the data.

Figure 4-3. Topology Diagram of a Distributed Switch That Handles VMkernel and Virtual Machine Networking

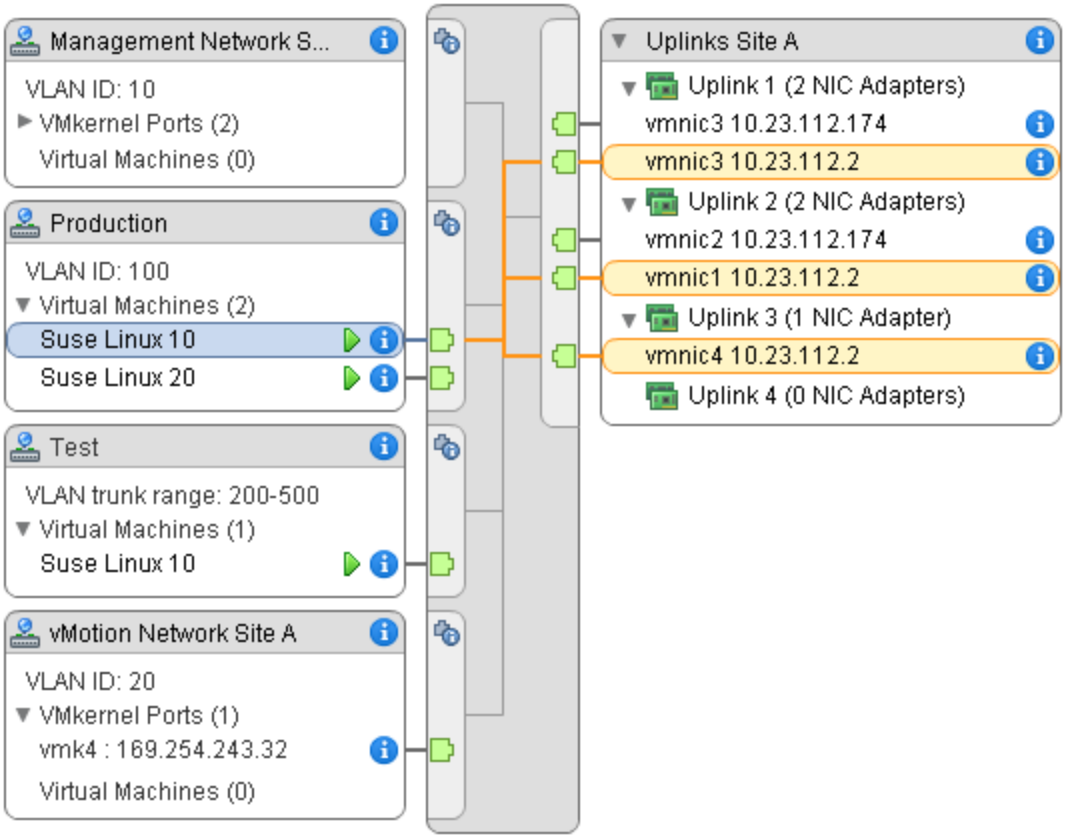

What to do next

You can perform the following common tasks in the topology of the distributed switch:

- Use filters to view the networking components only for selected port groups on certain hosts, for selected virtual machines, or for a port.
- Locate, configure and migrate virtual machine networking components across host and port groups by using the Migrate Virtual Machine Networking wizard.
- Detect the virtual machine adapters that have no network assigned and move them to the selected port group by using the Migrate Virtual Machine Networking wizard.
- Handle networking components on multiple hosts by using the Add and Manage Hosts wizard.
- View the physical NIC or NIC team that carries the traffic related to a selected virtual machine adapter or VMkernel adapter.

  In this way you can also view the host on which a selected VMkernel adapter resides. Select the adapter, trace the route to the associated physical NIC, and view the IP address or domain name next to the NIC.
- Determine the VLAN mode and ID for a port group. For information about VLAN modes, see "VLAN Configuration," on page 15.

View the Topology of a Host Proxy Switch in the vSphere Web Client

Examine and reorganize the networking of the VMkernel and virtual machines that the vSphere Distributed Switch handles on a host.

Procedure

1 In the vSphere Web Client, navigate to the host.

2 On the **Manage** tab, click **Networking**, and select **Virtual switches**.

3 Select the distributed switch from the list.

The topology of the host proxy switch appears under the list.

vSphere Distributed Switch Health Check

vSphere 5.1 distributed switch health check helps identify and troubleshoot configuration errors in vSphere distributed switches.

The following errors are common configuration errors that health check helps identify.

- Mismatched VLAN trunks between a vSphere distributed switch and physical switch.
- Mismatched MTU settings between physical network adapters, distributed switches, and physical switch ports.
- Mismatched virtual switch teaming policies for the physical switch port-channel settings.

Health check monitors the following:

- VLAN. Checks whether vSphere distributed switch VLAN settings match trunk port configuration on the adjacent physical switch ports.
- MTU. Checks whether the physical access switch port MTU jumbo frame setting based on per VLAN matches the vSphere distributed switch MTU setting.
- Teaming policies. Checks whether the physical access switch ports EtherChannel setting matches the distributed switch distributed port group IPHash teaming policy settings.

Health check is limited to only the access switch port to which the distributed switch uplink connects.

NOTE For VLAN and MTU checks, you must have at least two link-up physical uplink NICs for the distributed switch.

For a teaming policy check, you must have at least two link-up physical uplink NICs and two hosts when applying the policy.

Enable or Disable vSphere Distributed Switch Health Check in the vSphere Web Client

Health check monitors for changes in vSphere distributed switch configurations. You must enable vSphere distributed switch health check to perform checks on distributed switch configurations.

Health check is available only on ESXi 5.1 distributed switches. You can view health check information only through the vSphere Web Client 5.1 or later.

Procedure

1 Browse to a vSphere distributed switch in the vSphere Web Client.

2 Click the **Manage** tab.

3 Select **Settings**, and select **Health check**.

4 To enable or disable health check, click **Edit**.

5 Use the drop-down menus to enable or disable health check options.

Option	Description
VLAN and MTU	Reports the status of distributed uplink ports and VLAN ranges.
Teaming and Failover	Checks for any configuration mismatch between ESXi and the physical switch used in the teaming policy.

6 Click **OK**.

What to do next

When you change the configuration of a vSphere distributed switch, you can view information about the change in the **Monitor** tab in the vSphere Web Client. See "View vSphere Distributed Switch Health Check Information," on page 52.

View vSphere Distributed Switch Health Check Information

Once you have enabled health check, you can view vSphere distributed switch health check information in the vSphere Web Client.

Prerequisites

Enable health check on each vSphere distributed switch. See "Enable or Disable vSphere Distributed Switch Health Check in the vSphere Web Client," on page 51.

Procedure

1 Browse to a vSphere distributed switch in the vSphere Web Client.

2 Click the **Monitor** tab and click **Health**.

3 In the Health Status Details section, click a tab to view the health status of the selected check.

The three tabs include: **VLAN**, **MTU**, and **Teaming and Failover**.

Export, Import, and Restore Distributed Switch Configurations

You can export the configuration of a vSphere distributed switch from the vSphere Web Client, including distributed port group configurations. The exported configuration preserves valid networking settings, enabling their distribution to other deployments.

You can import or export a configuration of a distributed switch including its port groups. For information about exporting, importing, and restoring a port group configuration, see "Export, Import, and Restore vSphere Distributed Port Group Configurations," on page 44.

NOTE You can use a saved configuration file to restore policies and hosts associations on the distributed switch. You cannot restore the connection of physical NICs to uplink ports or ports of link aggregation groups.

Export vSphere Distributed Switch Configurations with the vSphere Web Client

You can export vSphere distributed switch and distributed port group configurations to a file. The file preserves valid network configurations, enabling distribution of these configurations to other deployments

This functionality is available only with the vSphere Web Client 5.1 or later. However, you can export settings from any version of a distributed switch if you use the vSphere Web Client or later.

Procedure

1 Browse to a distributed switch in the vSphere Web Client navigator.

2 Right-click the distributed switch in the navigator and select **All vCenter Actions > Export Configuration**.

3 Choose to export the distributed switch configuration, or export the distributed switch configuration and all port groups.

4 (Optional) Enter notes about this configuration in the **Descriptions** field.

5 Click **OK.**

6 Click **Yes** to save the configuration file to your local system.

You now have a configuration file that contains all the settings for the selected distributed switch and distributes port group. You can use this file to create multiple copies of this configuration on an existing deployment, or overwrite settings of existing distributed switches and port groups to conform to the selected settings.

What to do next

Use the exported configuration file to do the following tasks:

- To create a copy of the exported distributed switch, see "Import a vSphere Distributed Switch Configuration by Using the vSphere Web Client," on page 53.
- To overwrite settings on an existing distributed switch, see "Restore a vSphere Distributed Switch Configuration with the vSphere Web Client," on page 54.

You can also export, import, and restore only port group configurations. See "Export, Import, and Restore vSphere Distributed Port Group Configurations," on page 44.

Import a vSphere Distributed Switch Configuration by Using the vSphere Web Client

Import a stored configuration file to create a new vSphere Distributed Switch or to restore a switch that has been deleted earlier.

In vSphere 5.1 or later, you can import a distributed switch by using vSphere Web Client.

The configuration file contains the networking settings of the switch. By using it you can also replicate the switch in other virtual environments.

NOTE You can use a saved configuration file to replicate the switch instance, its host associations, and policies. You cannot replicate the connection of physical NICs to uplink ports or ports on link aggregation groups.

Procedure

1 In the vSphere Web Client, navigate to a datacenter.

2 Right-click the datacenter and select **All vCenter Actions > Import Distributed Switch**.

3 Browse to the location of the configuration file.

4 To assign the keys from the configuration file to the switch and its port groups, select the **Preserve original distributed switch and port group identifiers** check box and click **Next**.

If you select the **Preserve original distributed switch and port group identifiers** when you create a switch by using the configuration file of a deleted switch, all hosts that have been connected to the deleted switch are added again.

5 Review the settings for the switch and click **Finish**.

A new distributed switch is created with settings from the configuration file. If you have included distributed port group information in the configuration file, the port groups are also created.

Restore a vSphere Distributed Switch Configuration with the vSphere Web Client

Use the restore option to reset the configuration of an existing distributed switch to the settings in the configuration file. Restoring a distributed switch changes the settings on the selected switch back to the settings saved in the configuration file.

You can restore the configuration of a distributed switch only in vCenter Server and vSphere Web Client version 5.1 or later.

NOTE You can use a saved configuration file to restore policies and hosts associations on the distributed switch. You cannot restore the connection of physical NICs to uplink ports or ports of link aggregation groups.

Procedure

1 In the vSphere Web Client, navigate to the distributed switch.

2 Right-click the distributed switch in the navigator and select **All vCenter Actions > Restore Configuration**.

3 Browse for the configuration backup file to use.

4 Select **Restore distributed switch and all port groups** or **Restore distributed switch only** and click **Next**

5 Review the summary information for the restore.

Restoring a distributed switch will overwrite the current settings of the distributed switch and its port groups. It will not delete existing port groups that are not part of the configuration file.

6 Click **Finish**.

The distributed switch configuration has been restored to the settings in the configuration file.

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 in the vSphere Web Client

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

Procedure

1 Browse to a distributed switch in the vSphere Web Client navigator.

2 Click the **Manage** tab, and click **Settings**.

3 Select **Private VLAN** and click **Edit**.

4 Click **Add** to add a **Primary VLAN ID** to the list.

5 Click up and down arrows to select a primary private VLAN ID.

6 Click the **plus sign (+)** next to the Primary VLAN ID to add it to the list.

The primary private VLAN also appears under Secondary Private VLAN ID.

7 To add a secondary VLAN, click **Add** under the **Secondary VLAN** list, and click the up and down arrows to enter the number for the secondary VLAN.

8 Click the **plus sign (+)** next to the Secondary VLAN ID to add it to the list.

9 In the **Secondary VLAN type** column, click into the column to activate a drop-down menu. Select either **Isolated** or **Community** for the VLAN type.

10 Click **OK**.

What to do next

Configure a distributed port group or port to associate traffic with the private VLAN. See "Edit the VLAN Policy on a Distributed Port Group in the vSphere Web Client," on page 86 and "Edit the VLAN Policy on a Distributed Port with the vSphere Web Client," on page 87.

Remove a Primary Private VLAN with the vSphere Web Client

Remove unused primary private VLANs from the distributed settings view of the vSphere Web Client.

Prerequisites

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

Procedure

1 Browse to a distributed switch in the vSphere Web Client navigator.

2 Click the **Manage** tab, and click **Settings**.

3 Select **Private VLAN** and click **Edit**.

4 Select a primary private VLAN to remove.

When you remove a primary private VLAN, you also remove all associated secondary private VLANs.

5 Click **Remove** under the Primary VLAN ID list.

6 Click **OK** to verify that you want to remove the primary VLAN.

7 Click **OK.**

Remove a Secondary Private VLAN with the vSphere Web Client

Remove unused secondary private VLANs from the distributed settings view of the vSphere Web Client.

Prerequisites

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

Procedure

1 Browse to a distributed switch in the vSphere Web Client navigator.

2 Click the **Manage** tab, and select **Settings**.

3 Select **Private VLAN** and click **Edit**.

4 Select a primary private VLAN to display its associated secondary private VLANs.

5 Select a secondary private VLAN to remove.

6 Click **Remove** under the Secondary VLAN ID list, and click **OK**.

LACP Support on a vSphere Distributed Switch

The LACP support on a vSphere Distributed Switch 5.5 lets you connect ESXi hosts to physical switches by using dynamic link aggregation. You can create multiple link aggregation groups (LAGs) on a distributed switch to aggregate the bandwidth of physical NICs on ESXi hosts that are connected to LACP port channels.

Figure 4-4. Enhanced LACP Support on a vSphere Distributed Switch

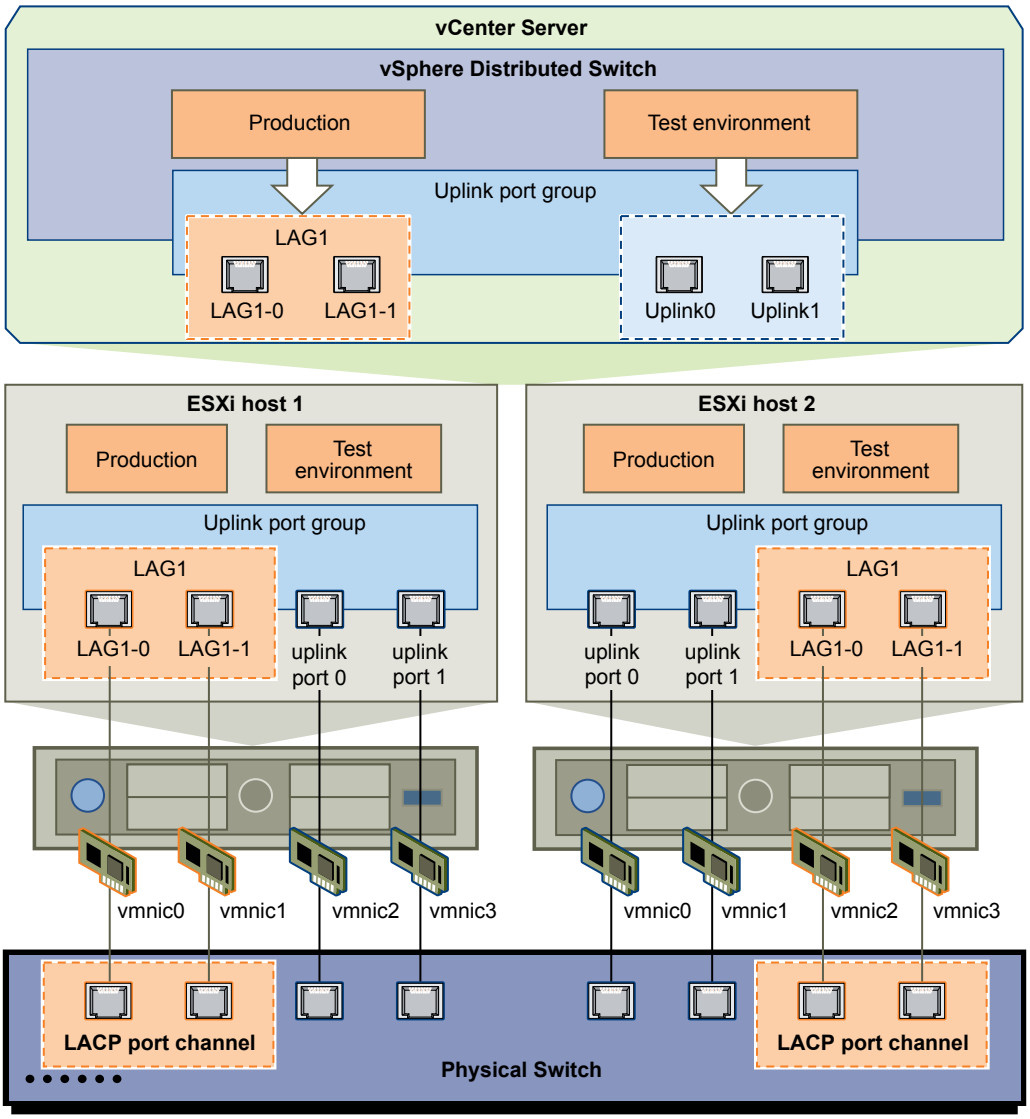

LACP Configuration on the Distributed Switch

You configure a LAG with two or more ports and connect physical NICs to the ports. LAG ports are teamed within the LAG, and the network traffic is load balanced between the ports through an LACP hashing algorithm. You can use a LAG to handle the traffic of distributed port groups to provide increased network bandwidth, redundancy, and load balancing to the port groups.

When you create a LAG on a distributed switch, a LAG object is also created on the proxy switch of every host that is connected to the distributed switch. For example, if you create LAG1 with two ports, LAG1 with the same number of ports is created on every host that is connected to the distributed switch.

On a host proxy switch, you can connect one physical NIC to only one LAG port. On the distributed switch, one LAG port can have multiple physical NICs from different hosts connected to it. The physical NICs on a host that you connect to the LAG ports must be connected to links that participate in an LACP port channel on the physical switch.

You can create up to 64 LAGs on a distributed switch. A host can support up to 32 LAGs. However, the number of LAGs that you can actually use depends on the capabilities of the underlying physical environment and the topology of the virtual network. For example, if the physical switch supports up to four ports in an LACP port channel, you can connect up to four physical NICs per host to a LAG.

Port Channel Configuration on the Physical Switch

For each host where you want to use LACP, you must create a separate LACP port channel on the physical switch. You must consider the following requirements when configuring LACP on the physical switch:

- The number of ports in the LACP port channel must be equal to the number of physical NICs that you want to group on the host. For example, if you want to aggregate the bandwidth of two physical NICs on a host, you must create an LACP port channel with two ports on the physical switch. The LAG on the distributed switch must be configured with at least two ports.
- The hashing algorithm of the LACP port channel on the physical switch must be the same as the hashing algorithm that is used for the LAG on the distributed switch.
- All the physical NICs that you want to connect to the LACP port channel must be configured with the same speed and duplex.

Comparison Between LACP Support in vSphere Distributed Switch 5.5 and 5.1

The LACP support on vSphere Distributed Switch 5.5 provides enhanced capabilities for managing link aggregation compared to vSphere Distributed Switch 5.1.

Table 4-1. Differences Between the LACP Support in vSphere Distributed Switch 5.1 and 5.5

Functionality	vSphere Distributed Switch 5.1	vSphere Distributed Switch 5.5	Description
Support for multiple LAGs	No	Yes	On vSphere Distributed Switch 5.1, the LACP support is enabled on an entire uplink port group and that port group acts as a single LAG for the switch. vSphere Distributed Switch 5.5 supports multiple LAGs.
Configure distributed port groups to use LAGs as active uplinks	No	Yes	On vSphere Distributed Switch 5.1, you can configure one LAG to handle the traffic for all distributed port groups on the distributed switch. vSphere Distributed Switch 5.5 lets you use a LAG to handle the traffic for individual distributed port groups. You can set LAGs as the active uplinks in the teaming and failover order of port groups.
Multiple LACP load balancing algorithms	No	Yes	The LACP support in vSphere Distributed Switch 5.1 supports only IP hash load balancing. In vSphere Distributed Switch 5.5, all load balancing algorithms of LACP are supported.

Convert to the Enhanced LACP Support on a vSphere Distributed Switch in the vSphere Web Client

After upgrading a vSphere Distributed Switch to version 5.5, you can convert to the enhanced LACP support to create multiple LAGs on the distributed switch.

If an LACP configuration exists on the distributed switch, enhancing the LACP support creates a new LAG and migrates all physical NICs from the standalone uplinks to the LAG ports. To create a different LACP configuration, you should disable the LACP support on the uplink port group before you start the conversion.

If the conversion to the enhanced LACP support fails, see *vSphere Troubleshooting* for details about how to complete it manually.

Prerequisites

- Verify that the vSphere Distributed Switch is version 5.5.
- Verify that none of the distributed port groups permit overriding their uplink teaming policy on individual ports.
- If you convert from an existing LACP configuration, verify that only one uplink port group exists on the distributed switch.
- Verify that hosts that participate in the distributed switch are connected and responding.
- Verify that you have the **dvPort group.Modify** privilege on the distributed port groups on the switch.
- Verify that you have the **Host.Configuration.Modify** privilege on the hosts on the distributed switch.

Procedure

1 In the vSphere Web Client, navigate to the distributed switch.

2 Select **Summary**.

3 In the Features section, click **Enhance** next to Link Aggregation Control Protocol.

4 (Optional) Select **Export configuration** to back up the configuration of the distributed switch and click **Next**.

The backup only stores the distributed switch configuration on the vCenter Server side. If the conversion to the enhanced LACP support fails, you can either use the backup to create a new distributed switch with the same configuration, or complete the conversion manually.

5 Review the validation prerequisites.

Prerequisite	Description
Port group accessibility	You have enough privileges to access and modify the uplink and distributed port groups on the switch.
LACP configuration	You have only one uplink port group on the distributed switch.
Uplink teaming policy override	Distributed port groups do not permit the override of their uplink teaming policy on individual ports.
Host accessibility	You have enough privileges to modify the networking configuration of the hosts connected to the distributed switch.
Host connectivity	Hosts that participate in the distributed switch are connected and responding.

6 Click **Next**.

7 If you convert from an existing LACP configuration, type the name of the new LAG in the Name text field.

8 Click **Next** to review the details about the conversion and click **Finish**.

You converted to the Enhanced LACP support on the vSphere Distributed Switch.

What to do next

Create LAGs on the distributed switch to aggregate the bandwidth of multiple physical NICs on the associated hosts.

LACP Teaming and Failover Configuration for Distributed Port Groups

To handle the network traffic of distributed port groups by using a LAG, you assign physical NICs to the LAG ports and set the LAG as active in the teaming and failover order of distributed port groups.

Table 4-2. LACP Teaming and failover configuration of distributed port groups

Failover Order	Uplinks	Description
Active	A single LAG	You can only use one active LAG or multiple standalone uplinks to handle the traffic of distributed port groups . You cannot configure multiple active LAGs or mix active LAGs and standalone uplinks.
Standby	Empty	Having an active LAG and standby uplinks and the reverse is not supported. Having a LAG and another standby LAG is not supported.
Unused	All standalone uplinks and other LAGs if any	Because only one LAG must be active and the Standby list must be empty, you must set all standalone uplinks and other LAGs to unused.

Configure a LAG to Handle the Traffic for Distributed Port Groups in the vSphere Web Client

To aggregate the bandwidth of multiple physical NICs on hosts, you can create a LAG on the distributed switch and use it to handle the traffic of distributed port groups.

Newly created LAGs do not have physical NICs assigned to their ports and are unused in the teaming and failover order of distributed port groups. To handle the network traffic of distributed port groups by using a LAG, you must migrate the traffic from standalone uplinks to the LAG.

Prerequisites

- Verify that for every host where you want to use LACP, a separate LACP port channel exists on the physical switch. See "LACP Support on a vSphere Distributed Switch," on page 57.
- Verify that the vSphere Distributed Switch where you configure the LAG is version 5.5.
- Verify that enhanced LACP is supported on the distributed switch.

Procedure

1 Create a LAG in the vSphere Web Client on page 61

To migrate the network traffic of distributed port groups to a LAG, you create a new LAG.

2 Set a LAG as Standby in the Teaming and Failover Order of Distributed Port Groups in the vSphere Web Client on page 62

The new LAG by default is unused in the teaming and failover order of distributed port groups. Because only one LAG or only standalone uplinks can be active for distributed port groups, you must create an intermediate teaming and failover configuration, where the LAG is standby. This configuration lets you migrate physical NICs to the LAG ports by keeping the network connectivity up.

3 Assign Physical NICs to LAG Ports in the vSphere Web Client on page 62

You have set the new LAG as standby in the teaming and failover order of distributed port groups. Having the LAG as standby lets you safely migrate the physical NICs from standalone uplinks to the LAG ports without losing network connectivity.

4 Set the LAG as Active in the Teaming and Failover Order of the Distributed Port Group in the vSphere Web Client on page 63

You migrated physical NICs to the LAG ports. Set the LAG as active and move all standalone uplinks as unused in the teaming and failover order of the distributed port groups.

Create a LAG in the vSphere Web Client

To migrate the network traffic of distributed port groups to a LAG, you create a new LAG.

Procedure

1 In the vSphere Web Client, navigate to the distributed switch.

2 Select **Manage**, and select **Settings**.

3 Under **LACP**, click **New Link Aggregation Group**.

4 Name the new LAG.

5 Set the number of ports to the LAG.

Set the same number of ports to the LAG as the number of ports in the LACP port channel on the physical switch. A LAG port has the same function as an uplink on the distributed switch. All LAG ports form a NIC team in the context of the LAG.

6 Select the LACP negotiating mode of the LAG.

Option	Description
Active	All LAG ports are in an Active negotiating mode. The LAG ports initiate negotiations with the LACP port channel on the physical switch by sending LACP packets.
Passive	The LAG ports are in Passive negotiating mode. They respond to LACP packets they receive but do not initiate LACP negotiation.

If the LACP-enabled ports on the physical switch are in Active negotiating mode, you can set the LAG ports in Passive mode and the reverse.

7 Select a load balancing mode from the hashing algorithms that LACP defines.

NOTE The hashing algorithm must be the same as the hashing algorithm set to the LACP port channel on the physical switch.

8 Set the VLAN and the NetFlow policies for the LAG.

This option is active when overriding the VLAN and NetFlow policies per individual uplink ports is enabled on the uplink port group. If you set the VLAN and NetFlow policies to the LAG, they override the policies set on the uplink port group level.

9 Click **OK**.

The new LAG is unused in the teaming and failover order of distributed port groups. No physical NICs are assigned to the LAG ports.

As with standalone uplinks, the LAG has a representation on every host that is associated with the distributed switch. For example, if you create LAG1 with two ports on the distributed switch, a LAG1 with two ports is created on every host that is associated with the distributed switch.

What to do next

Set the LAG as standby in the teaming and failover configuration of distributed port groups. In this way, you create an intermediate configuration that lets you migrate the network traffic to the LAG without losing network connectivity.

Set a LAG as Standby in the Teaming and Failover Order of Distributed Port Groups in the vSphere Web Client

The new LAG by default is unused in the teaming and failover order of distributed port groups. Because only one LAG or only standalone uplinks can be active for distributed port groups, you must create an intermediate teaming and failover configuration, where the LAG is standby. This configuration lets you migrate physical NICs to the LAG ports by keeping the network connectivity up.

Procedure

1 Navigate to the distributed switch.

2 From the **Actions** menu, select **Manage Distributed Port Groups**.

3 Select **Teaming and failover** and click **Next**.

4 Select the port groups where you want to use the LAG.

5 In Failover order, select the LAG and use the up arrow to move it to the Standby uplinks list.

6 Click **Next**, review the message that informs you about the usage of the intermediate teaming and failover configuration, and click **OK**.

7 On the Ready to complete page, click **Finish**.

What to do next

Migrate physical NICs from standalone uplinks to the LAG ports.

Assign Physical NICs to LAG Ports in the vSphere Web Client

You have set the new LAG as standby in the teaming and failover order of distributed port groups. Having the LAG as standby lets you safely migrate the physical NICs from standalone uplinks to the LAG ports without losing network connectivity.

Prerequisites

- Verify that either all LAG ports or the corresponding LACP-enabled ports on the physical switch are in active LACP negotiating mode.
- Verify that the physical NICs that you want to assign to the LAG ports have the same speed and are configured at full duplex.

Procedure

1 In the vSphere Web Client, navigate to the distributed switch where the LAG resides.

2 From the **Actions** menu, select **Add and Manage Hosts**.

3 Select **Manage host networking**.

4 Select the host whose physical NICs you want to assign to the LAG ports and click **Next**.

5 On the Select network adapter tasks page, select **Manage physical adapters** and click **Next**.

6 On the Manage physical network adapters page, select a NIC and click **Assign an uplink**.

7 Select a LAG port and click **OK**.

8 Repeat Step 6 and Step 7 for all physical NICs that you want to assign to the LAG ports.

9 Complete the wizard.

Example: Configure Two Physical NICs to a LAG in the Add and Manage Hosts Wizard

For example, if you have a LAG with two ports, you configure a physical NIC to each LAG port in the Add and Manage Hosts wizard.

Figure 4-5. Connecting Two Physical NICs to LAG Ports

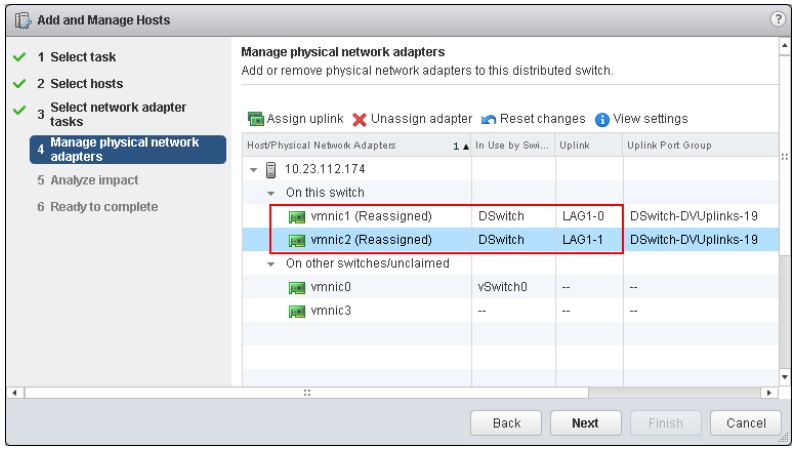

What to do next

Set the LAG as active and all standalone uplinks to unused in the teaming and failover order of distributed port groups.

Set the LAG as Active in the Teaming and Failover Order of the Distributed Port Group in the vSphere Web Client

You migrated physical NICs to the LAG ports. Set the LAG as active and move all standalone uplinks as unused in the teaming and failover order of the distributed port groups.

Procedure

1 Navigate to the distributed switch.

2 From the **Actions** menu, select **Manage Distributed Port Groups**.

3 Select **Teaming and failover** and click **Next**.

4 Select the port groups where you set the LAG as standby and click **Next**.

5 In Failover order, use the up and down arrows to move the LAG in the Active list, all standalone uplinks in the Unused list, and leave the Standby list empty.

6 Click **Next** and click **Finish**.

You safely migrated network traffic from standalone uplinks to a LAG for distributed port groups and created a valid LACP teaming and failover configuration for the groups.

Example: Topology of a Distributed Switch that Uses a LAG

If you configure a LAG with two ports to handle the traffic of a distributed port group, you can check the topology of the distributed switch to view how it changed as a result of the new configuration.

Figure 4-6. Distributed Switch Topology with a LAG

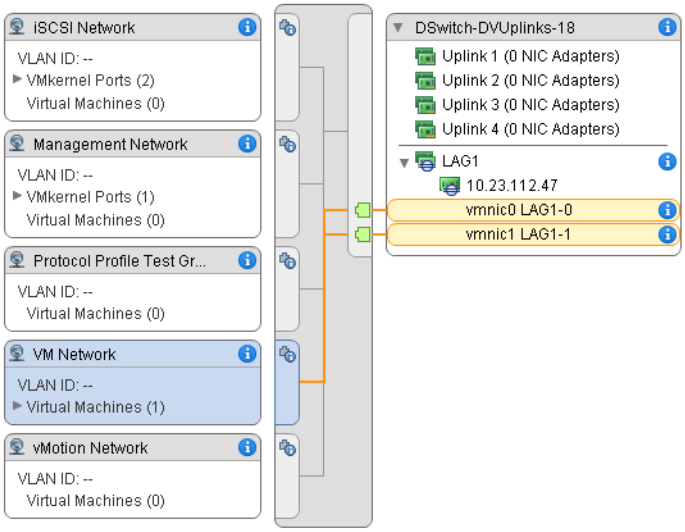

Edit a LAG in the vSphere Web Client

Edit the settings of a LAG if you need to add more ports to the group or change the LACP negotiating mode, the load balancing algorithm, or the VLAN and NetFlow policies.

Procedure

1 In the vSphere Web Client, navigate to the vSphere Distributed Switch.

2 Select **Manage** and select **Settings**.

3 Select **LACP**.

4 In the **Name** text box, type a new name for the LAG.

5 Change the number of ports for the LAG if you want to add more physical NICs to it.

The new NICs must be connected to ports that are part of an LACP port channel on the physical switch.

6 Change the LACP negotiating mode of the LAG.

If all ports on the physical LACP port channel are in Active LACP mode, you can change the LACP mode of the LAG to Passive and the reverse.

7 Change the load balancing mode of the LAG.

You can select from the load balancing algorithms that LACP defines.

8 Change the VLAN and the NetFlow policies.

This option is active when the option for overriding the VLAN and NetFlow policies for individual ports is enabled on the uplink port group. If you change the VLAN and NetFlow policies for the LAG, they override the policies set at the uplink port group level.

9 Click **OK**.

Enable LACP 5.1 Support on an Uplink Port Group in the vSphere Web Client

You can enable LACP support on an uplink port group for vSphere Distributed Switches 5.1, and for switches upgraded to 5.5 that do not have the enhanced LACP support.

Prerequisites

- Verify that for each host where you want to use LACP, a separate LACP port channel exists on the physical switch.
- Verify that distributed port groups have their load balancing policy set to IP hash.
- Verify that the LACP port channel on the physical switch is configured with IP hash load balancing.
- Verify that distributed port groups have network failure detection policy set to link status only.
- Verify that distributed port groups have all uplinks set to active in their teaming and failover order.
- Verify that all physical NICs that are connected to the uplinks have the same speed and are configured at full duplex.

Procedure

1 In the vSphere Web Client, navigate to an uplink port group.

   a Select a distributed switch and click **Related Objects**.

   b Click **Uplink Port Groups** and select the uplink port group.

2 Click **Manage**, and click **Settings**.

3 Click **Edit**.

4 In the LACP section, use the drop-down list to enable LACP.

5 Set the LACP negotiating mode for the uplink port group.

Option	Description
Active	All uplink ports in the group are in an Active negotiating mode. The uplink ports initiate negotiations with the LACP-enabled ports on the physical switch by sending LACP packets.
Passive	All uplink ports are in a Passive negotiating mode. They respond to LACP packets that they receive but do not initiate LACP negotiation.

   If the LACP-enabled ports on the physical switch are in Active negotiating mode, you can set the uplink ports in Passive mode and the reverse.

6 Click **OK**.

Limitations of the LACP Support on a vSphere Distributed Switch

The LACP support on a vSphere Distributed Switch lets network devices to negotiate automatic bundling of links by sending LACP packets to a peer. However, the LACP support on a vSphere Distributed Switch has limitations.

- The LACP support is not compatible with software iSCSI multipathing.
- The LACP support settings are not available in host profiles.
- The LACP support is not possible between nested ESXi hosts.
- The LACP support does not work with the ESXi dump collector.
- The LACP support does not work with port mirroring.

- The teaming and failover health check does not work for LAG ports. LACP checks the connectivity of the LAG ports.
- The enhanced LACP support works correctly when only one LAG handles the traffic per distributed port or port group.
- The LACP 5.1 support only works with IP Hash load balancing and Link Status Network failover detection.
- The LACP 5.1 support only provides one LAG per distributed switch and per host.

Setting Up VMkernel Networking 5

You set up VMkernel adapters to provide network connectivity to hosts and to accommodate the traffic of vMotion, IP storage, Fault Tolerance logging, and Virtual SAN.

- Remove a VMkernel Adapter in the vSphere Web Client on page 75

  Remove a VMkernel adapter from a vSphere distributed or a standard switch when you no longer need the adapter. Make sure that you leave at least one VMkernel adapter for management traffic on the host to keep the network connectivity up.

The VMkernel Networking Layer

The VMkernel networking layer provides connectivity to hosts and handles the standard infrastructure traffic of vSphere vMotion, IP storage, Fault Tolerance, and Virtual SAN. You can set up VMkernel adapters for the standard infrastructure traffic on vSphere standard switches and on vSphere distributed switches.

TCP/IP Stacks at the VMkernel Level

The default TCP/IP stack at the VMkernel level provides networking support for the standard infrastructure types of traffic. You can add custom TCP/IP stacks at the VMkernel level and forward networking traffic through custom applications.

Securing Infrastructure Traffic

You should take appropriate security measures to prevent unauthorized access to the management and application traffic in your vSphere environment. For example, you should isolate the vMotion traffic in a separate network that includes only the ESXi hosts that participate in the migration. You should isolate the management traffic in a network that only network and security administrators are able to access. For more information, see *vSphere Security* and *vSphere Installation and Setup*.

Infrastructure Traffic Types on the Default TCP/IP Stack

You should dedicate a separate VMkernel adapter for every traffic type. For distributed switches, dedicate a separate distributed port group for each VMkernel adapter.

Management traffic

It carries the configuration and management communication for ESXi hosts and vCenter Server as well as the host-to-host High Availability traffic. By default, when you install the ESXi software, a vSphere Standard switch is created on the host together with a VMkernel adapter for management traffic. To provide redundancy and increase bandwidth, you can connect two or more physical NICs to a VMkernel adapter for management traffic.

vMotion traffic

To accommodate the vMotion traffic, a VMkernel adapter for vMotion is required both on the source and the target hosts. The VMkernel adapters for vMotion should handle only the vMotion traffic. For better performance, you can assign multiple physical NICs to the port group of the VMkernel adapter. In this way, multiple physical NICs are used for vMotion, which results in greater bandwidth.

NOTE vMotion network traffic is not encrypted. You should provision secure private networks for use by vMotion only.

IP storage traffic

Storage types that use standard TCP/IP networks and depend on the VMkernel networking layer require VMkernel adapters. Such storage types are software iSCSI, depended hardware iSCSI, and NFS. If you have two or more physical NICs for iSCSI, you can configure iSCSI multipathing. NFS does not require a separate VMkernel adapter. It uses the management traffic on the host for I/O. ESXi hosts support only NFS version 3 over TCP/IP.

Fault Tolerance traffic The traffic that the primary fault tolerant virtual machine sends to the secondary fault tolerant virtual machine over the VMkernel networking layer. A separate VMkernel adapter for Fault Tolerance logging is required on every host that is part of a vSphere HA cluster.

Virtual SAN traffic Every host that participates in a Virtual SAN cluster must have a VMkernel adapter to handle the Virtual SAN traffic.

View Information About VMkernel Adapters on a Host in the vSphere Web Client

For each VMkernel adapter, you can view its assigned services, associated switch, port settings, IP settings, TCP/IP stack, VLAN ID, and policies.

Procedure

1 In the vSphere Web Client, navigate to the host.

2 Click **Manage** and click **Networking**.

3 To view information about all VMkernel adapters on the host, select **VMkernel adapters**.

For ESX hosts 4.x, the VMkernel adapters option displays the ESX service console.

Option	Description
Device	Name of the VMkernel adapter.
Network Label	Name of the network to which the VMkernel adapter is connected.
Switch	vSphere standard or distributed switch with which the VMkernel adapter is associated.
IP Address	The IP address of the VMkernel adapter.
TCP/IP Stack	The TCP/IP stack that handles the traffic for the VMkernel adapter. If there is a custom TCP/IP stack, set for the VMkernel adapter, it is displayed.
vMotion Traffic	Status of vMotion on the VMkernel adapter.
FT Logging	Status of FT Logging on the VMkernel adapter.
Management Traffic	Status of Management Traffic on the VMkernel adapter.
Virtual SAN Traffic	Status of the Virtual SAN traffic on the VMkernel adapter.

4 Select an adapter from the VMkernel adapters list to view its settings.

Tab	Description
All	Displays all configuration information about the VMkernel adapter. This includes port and NIC settings, IPv4 and IPv6 settings, traffic shaping, teaming and failover, and security policies.
Properties	Displays the port properties and NIC settings of the VMkernel adapter. The port properties include the port group (network label) to which the adapter is associated, the VLAN ID and the enabled services. The NIC settings include
IP Settings	Displays all IPv4 and IPv6 settings for the VMkernel adapter. IPv6 information is not displayed if IPv6 has not been enabled on the host.
Policies	Displays the configured traffic shaping, teaming and failover, and security policies that apply for the port group where the VMkernel adapter is connected.

Create a VMkernel Adapter on a vSphere Standard Switch in the vSphere Web Client

Create a VMkernel network adapter on a vSphere standard switch to provide network connectivity for hosts and to handle the traffic for vSphere vMotion, IP storage, Fault Tolerance logging, or Virtual SAN. You should dedicate a VMkernel adapter to only one traffic type.

Procedure

1 In the vSphere Web Client, navigate to the host.

2 Under **Manage**, select **Networking** and then select **VMkernel adapters**.

3 Click **Add host networking**.

4 On the Select connection type page, select **VMkernel Network Adapter** and click **Next**.

5 On the Select target device page, select either an existing standard switch or a **New vSphere standard switch**.

6 (Optional) On the Create a Standard Switch page, assign physical NICs to the switch.

You can create the standard switch without physical NICs and configure them later. During that time the host does not have network connectivity to the other hosts on the physical network. The virtual machines on the host are able to communicate with each other.

a Click **Add** and select as many physical NICs as you need.

b Use the up and down arrows to configure the active and standby NICs.

7 On the Port properties page, configure the settings for the VMkernel adapter.

Option	Description
Network label	Type a value for this label to indicate the traffic type for the VMkernel adapter, for example **`Management traffic`** or **`vMotion`**.
VLAN ID	Set a VLAN ID to identify the VLAN that the network traffic of the VMkernel adapter will use.
IP settings	Select IPv4, IPv6, or both. NOTE The IPv6 option does not appear on hosts that do not have IPv6 enabled.
TCP/IP stack	Select a stack from the list.
Enable services	You can enable services for the default TCP/IP stack on the host. Select from the available services: ■ **vMotion traffic**. Enables the VMkernel adapter to advertise itself to another host as the network connection where vMotion traffic is sent. You can enable this property for only one vMotion and IP storage VMkernel adapter per host. If this property is not enabled for any VMkernel adapter, migration with vMotion to the selected host is not possible. ■ **Fault Tolerance traffic**. Enables Fault Tolerance logging on the host. ■ **Management traffic**. Enables the management traffic for the host and vCenter Server. Typically, hosts have such a VMkernel adapter created when the ESXi software was installed. You can create another VMkernel adapter for management traffic on the host to provide redundancy. ■ **Virtual SAN**. Enables the Virtual SAN traffic on the host. Every host that is part from a Virtual SAN cluster must have such a VMkernel adapter.

8 (Optional) On the IPv4 settings page, select an option for obtaining IP addresses.

Option	Description
Obtain IP settings automatically	Use DHCP to obtain IP settings.
Use static IP settings	Type the IPv4 IP address and subnet mask for the VMkernel adapter. The VMkernel Default Gateway and DNS server addresses for IPv4 are obtained from the selected TCP/IP stack.

9 (Optional) On the IPv6 settings page, select an option for obtaining IPv6 addresses.

Option	Description
Obtain IPv6 addresses automatically through DHCP	Use DHCP to obtain IPv6 addresses.
Obtain IPv6 addresses automatically through Router Advertisement	Use router advertisement to obtain IPv6 addresses.
Static IPv6 addresses	a Click **Add** to add a new IPv6 address. b Type the IPv6 address and subnet prefix length, and click **OK**. c To change the VMkernel default gateway, click **Edit**.

10 Review your setting selections in the Ready to complete page and click **Finish**.

Create a VMkernel Adapter on a Host Associated with a vSphere Distributed Switch in the vSphere Web Client

Create a VMkernel adapter on a host that is associated with a distributed switch to provide network connectivity to the host and to handle the traffic for vSphere vMotion, IP storage, Fault Tolerance logging, and Virtual SAN. You should dedicate a single distributed port group per VMkernel adapter. One VMkernel adapter should handle one traffic type.

Procedure

1 In the vSphere Web Client, navigate to the host.

2 Under **Manage**, select **Networking** and then select **VMkernel adapters**.

3 Click **Add host networking**.

4 On the Select connection type page, select **VMkernel Network Adapter** and click **Next**.

5 Select a distributed port group and click **Next**.

6 On the Port properties page, configure the settings for the VMkernel adapter.

Option	Description
Network label	The network label is inherited from the label of the distributed port group.
IP settings	Select IPv4, IPv6, or both. NOTE The IPv6 option does not appear on hosts that do not have IPv6 enabled.

Option	Description
TCP/IP stack	If custom stacks are available, select one from the list.
Enable services	You can enable services for the default TCP/IP stack on the host. Select from the available services: ■ **vMotion traffic**. Enables the VMkernel adapter to advertise itself to another host as the network connection where vMotion traffic is sent. You can enable this property for only one vMotion and IP storage VMkernel adapter per host. If this property is not enabled for any VMkernel adapter, migration with vMotion to the selected host is not possible. ■ **Fault Tolerance traffic**. Enables Fault Tolerance logging on the host. ■ **Management traffic**. Enables the management traffic for the host and vCenter Server. Typically, hosts have such a VMkernel adapter created when the ESXi software is installed. You can create another VMkernel adapter for management traffic on the host to provide redundancy. ■ **Virtual SAN**. Enables the Virtual SAN traffic on the host. Every host that is part of a Virtual SAN cluster must have such a VMkernel adapter.

7 (Optional) On the IPv4 settings page, select an option for obtaining IP addresses.

Option	Description
Obtain IP settings automatically	Use DHCP to obtain IP settings.
Use static IP settings	Type the IPv4 IP address and subnet mask for the VMkernel adapter. The VMkernel Default Gateway and DNS server addresses for IPv4 are obtained from the selected TCP/IP stack.

8 (Optional) On the IPv6 settings page, select an option for obtaining IPv6 addresses.

Option	Description
Obtain IPv6 addresses automatically through DHCP	Use DHCP to obtain IPv6 addresses.
Obtain IPv6 addresses automatically through Router Advertisement	Use router advertisement to obtain IPv6 addresses.
Static IPv6 addresses	a Click **Add** to add a new IPv6 address. b Type the IPv6 address and subnet prefix length, and click **OK**. c To change the VMkernel default gateway, click **Edit**.

9 Review your setting selections in the Ready to complete page and click **Finish**.

Edit a VMkernel Adapter Configuration in the vSphere Web Client

You might need to change the supported traffic type for a VMkernel adapter, or the way IPv4 or IPv6 addresses are obtained.

Procedure

1 In the vSphere Web Client, navigate to the host.

2 Under **Manage**, select **Networking** and then select **VMkernel adapters**.

3 Select the VMkernel adapter that resides on the target distributed or standard switch and click **Edit**.

4 On the Port properties page, select the services that you want to enable.

Check box	Description
vMotion traffic	Enables the VMkernel adapter to advertise itself to another host as the network connection where vMotion traffic is sent. You can enable this property for only one vMotion and IP storage VMkernel adapter per host. If this property is not enabled for any VMkernel adapter, migration with vMotion to the selected host is not possible.
Fault Tolerance traffic	Enables Fault Tolerance logging on the host.
Management traffic	Enables the management traffic for the host and vCenter Server. Typically, hosts have such a VMkernel adapter created when the ESXi software was installed. You can have an additional VMkernel adapter for management traffic on the host to provide redundancy.
Virtual SAN	Enables Virtual SAN traffic on the host. Every host that is part from a Virtual SAN cluster must have such a VMkernel adapter.

5 On the NIC Settings page, set the MTU for the network adapter.

6 With IPv4 enabled, in the IPv4 settings section select the method by which IP addresses are obtained.

Option	Description
Obtain IP settings automatically	Use DHCP to obtain IP settings.
Use static IP settings	Type the IPv4 IP address and subnet mask for the VMkernel adapter. The VMkernel Default Gateway and DNS server addresses for IPv4 are obtained from the selected TCP/IP stack.

7 With IPv6 enabled, in the IPv6 settings select an option for obtaining IPv6 addresses.

NOTE The IPv6 option does not appear on hosts that do not have IPv6 enabled.

Option	Description
Obtain IPv6 addresses automatically through DHCP	Use DHCP to obtain IPv6 addresses.
Obtain IPv6 addresses automatically through Router Advertisement	Use router advertisement to obtain IPv6 addresses.
Static IPv6 addresses	a Click **Add** to add a new IPv6 address. b Type the IPv6 address and subnet prefix length, and click **OK**. c To change the VMkernel default gateway, click **Edit**.

In the Advanced Settings section of IP Settings, you can remove IPv6 addresses. If router advertisement is enabled, removed addresses from this origin might reappear. Removal of DHCP addresses on the VMkernel adapter is not supported. These addresses are removed only when the DHCP option is turned off.

8 On the Validate changes page, verify that the changes made to the VMKernel adapter will not disrupt other operations.

9 Click **OK**.

View TCP/IP Stack Configuration on a Host in the vSphere Web Client

You can view the DNS and routing configuration of a TCP/IP stack on a host, the IPv4 and IPv6 routing tables, the congestion control algorithm, and the maximum number of allowed connections.

Procedure

1 In the vSphere Web Client, navigate to the host.

2 Click **Manage**, click **Networking**, and select **TCP/IP configuration**.

3 Select a stack from the TCP/IP Stacks table.

If there are no custom TCP/IP stacks configured on the host, you view only the default TCP/IP stack on the host.

DNS and routing details about the selected TCP/IP stack appear below the TCP/IP Stacks table. You can view the IPv4 and IPv6 routing tables, and the DNS and routing configuration for the stack.

NOTE The IPv6 routing table is only visible if IPv6 is enabled on the host.

The **Advanced** tab contains information about the configured congestion control algorithm and the maximum number of allowed connections to the stack.

Change the Configuration of a TCP/IP Stack on a Host in the vSphere Web Client

You can change the DNS and default gateway configuration of a TCP/IP stack on a host, the congestion control algorithm, the maximum number of connections, and the name of custom TCP/IP stacks.

Procedure

1 In the vSphere Web Client, navigate to the host.

2 Click **Manage**, click **Networking**, and select **TCP/IP configuration**.

3 Select a stack from the table and click **Edit**.

4 On the Name page, you can change the name of a custom TCP/IP stack.

5 On the DNS configuration page, select a method of obtaining DNS server information.

Option	Description
Obtain settings automatically from virtual network adapter	From the **VMKernel network adapter** drop-down menu, select a network adapter.
Enter settings manually	a Edit the Host name. b Edit the Domain name. c Type a preferred DNS server IP address. d Type an alternate DNS server IP address. e (Optional) Use the **Search domains** text box to look for hosts with specific names.

6 On the Routing page, edit the VMkernel gateway information.

NOTE Removing the default gateway might cause the client to lose connectivity with the host.

7 On the Advanced page, edit the congestion control algorithm of the stack and the maximum number of connections.

8 Click **OK**.

Create a Custom TCP/IP Stack

You can create a custom TCP/IP stack on a host to forward VMkernel traffic through a custom application.

Procedure

1 Open an SSH connection to the host.

2 Log in as the root user.

3 Run the following vSphere CLI command:

```
esxcli network ip netstack add –N="stack_name"
```

The custom TCP/IP stack is created on the host. You can assign VMkernel adapters to the stack.

Remove a VMkernel Adapter in the vSphere Web Client

Remove a VMkernel adapter from a vSphere distributed or a standard switch when you no longer need the adapter. Make sure that you leave at least one VMkernel adapter for management traffic on the host to keep the network connectivity up.

Procedure

1 In the vSphere Web Client, navigate to the host.

2 Under **Manage**, select **Networking** and then select **VMkernel adapters**.

3 Select a VMkernel adapter from the list, and click **Remove**.

4 On the confirmation dialog, click **Analyze impact**.

5 Review the impacted services and the level of impact.

Option	Description
No impact	The service will continue its normal function after the new networking configuration is applied.
Important impact	The normal function of the service might be disrupted if the new networking configuration is applied.
Critical impact	The normal function of the service will be interrupted if the new networking configuration is applied.

   a If impact on a service is important or critical, click the service and review the reasons that are displayed in the Analysis details pane.

   b If there are no impacted services, close the Analyze Impact dialog, otherwise cancel the removal of the VMkernel adapter until you fix the seasons for any critical or important impact on a service.

6 Click **OK**.

Networking Policies 6

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.

- Teaming and Failover Policy on page 78

  Teaming and failover policy lets you to determine how network traffic is distributed between physical adapters and how to reroute traffic in the event of adapter failure.

- VLAN Policy on page 86

  VLAN policies determine how VLANs function across your network environment.

- Security Policy on page 88

  Networking security policy provides protection of traffic against MAC address impersonation and unwanted port scanning

- Traffic Shaping Policy on page 92

  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.

- Resource Allocation Policy on page 96

  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.

- Monitoring Policy on page 97

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

- Traffic Filtering and Marking Policy on page 98

  In a vSphere distributed switch 5.5 and later, by using the traffic filtering and marking policy, you can protect the virtual network from unwanted traffic and security attacks or apply a QoS tag to a certain type of traffic.

- Port Blocking Policies on page 114

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

- Manage Policies for Multiple Port Groups on a vSphere Distributed Switch in the vSphere Web Client on page 115

  You can modify networking policies for multiple port groups on a vSphere Distributed Switch.

Teaming and Failover Policy

Teaming and failover policy lets you to determine how network traffic is distributed between physical adapters and how to reroute traffic in the event of adapter failure.

You can edit your load balancing and failover policy by configuring the following parameters:

- Load Balancing 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.
- Failover Order can be active or standby.

You might lose connectivity when a failover event occurs. This loss causes the MAC addresses used by virtual machines that are associated with standard or distributed switch to appear on a different physical switch port than the one they had been on previously. To avoid this problem, put your physical switch in Port Fast or Port Fast trunk mode.

Edit Teaming and Failover Policy for a vSphere Standard Switch in the vSphere Web Client

Use the teaming and failover policy to determine how network traffic of the virtual machines and VMkernel adapters that are connected to the switch is distributed between physical adapters and how to reroute traffic in the event of an adapter failure.

Procedure

1 In the vSphere Web Client, navigate to the host.

2 On the **Manage** tab, click **Networking**, and select **Virtual switches**.

3 Select a standard switch from the list, click **Edit settings** and select **Teaming and Failover**.

4 From the **Load Balancing** drop-down menu, specify how the standard or distributed switch selects an uplink to handle traffic from a virtual machine or VMkernel adapter.

Option	Description
Route based on the originating virtual port	Select an uplink based on the virtual port where the traffic entered the virtual switch.
Route based on IP hash	Select an uplink based on a hash of the source and destination IP addresses of each packet. For non-IP packets, the switch simply uses the data at those fields to compute the hash. IP-based teaming requires that the physical switch be configured with EtherChannel.
Route based on source MAC hash	Select an uplink based on a hash of the source Ethernet.
Route based on physical NIC load	Available for distributed port groups or distributed ports. Select an uplink based on the current load of the physical network adapters connected to the port group or port. If an uplink remains busy at 75% or higher for 30 seconds, the host proxy switch moves a part of the virtual machine traffic to a physical adapter that has free capacity.
Use explicit failover order	From the list of active adapters, always use the highest order uplink that passes failover detection criteria.

5 From the **Network Failover Detection** drop-down menu, specify the method that the standard or distributed switch uses for failover detection.

Option	Description
Link Status only	Relies only on the link status that the network adapter provides. This option detects failures, such as removed cables and physical switch power failures. However, it does not detect configuration errors, such as the following: ■ Physical switch port that is blocked by spanning tree or is misconfigured to the wrong VLAN. ■ Pulled cable that connects a physical switch to another networking devices, for example, an upstream switch.
Beacon Probing	Sends out and listens for beacon probes on all NICs in the team and uses this information, in addition to link status, to determine link failure. ESX/ESXi sends beacon packets every second. Beaconing is most useful with three or more NICs in a team because ESX/ESXi can detect failures of a single adapter. If only two NICs are assigned and one of them loses connectivity, the switch cannot determine which NIC needs to be taken out of service because both do not receive beacons and as a result all packets sent to both uplinks. Using at least three NICs in such a team allows for $n-2$ failures where n is the number of NICs in the team before reaching an ambiguous situation. The NICs must be in an active/active or active/standby configuration because the NICs in an unused state do not participate in beacon probing.

6 From the **Notify Switches** drop-down menu, select whether the standard or distributed switch notifies the physical switch in case of a failover.

If you select **Yes**, whenever a virtual NIC is connected to the virtual 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 physical switches. Notifying the physical switch offers lowest latency when a failover or a migration with vSphere vMotion occurs.

NOTE Set this option to **No** if a connected virtual machines is using Microsoft Network Load Balancing in unicast mode. No issues exist with Network Load Balancing running in multicast mode.

7 From the **Failback** drop-down menu, determine whether a physical adapter is returned to active status after recovering from a failure.

If failback is set to **Yes** (default), the adapter is returned to active duty immediately upon recovery, displacing the standby adapter that took over its slot, if any.

If failback is set to **No** to a standard port, a failed adapter is left inactive after recovery until another currently active adapter fails and must be replaced.

8 Specify how the uplinks in a team are used when a failover occurs by configuring the Failover Order list.

If you want to use some uplinks but reserve others for emergencies in case the uplinks in use fail, use the up and down arrow keys to move them into different groups.

Option	Description
Active adapters	Continue to use the uplink if the network adapter connectivity is up and active.
Standby adapters	Use this uplink if one of the active physical adapter is down.
Unused adapters	Do not use this uplink.

9 Click **OK**.

Edit the Teaming and Failover Policy on a Standard Port Group in the vSphere Web Client

Configure the teaming and failover policy on a standard port group to determine how network traffic related to a group of virtual machines or a VMkernel adapter is distributed between physical adapters and how to reroute traffic in the event of a physical adapter failure.

Procedure

1 In the vSphere Web Client, navigate to the host.

2 On the **Manage** tab, click **Networking**, and select **Virtual switches**.

3 Select a standard switch from the list.

The topology diagram of the switch appears.

4 In the topology diagram, select the port group and click **Edit settings**.

5 On the Teaming and Failover page, to override the teaming and failover properties inherited from the standard switch, select the check boxes next to the properties that you want to override.

6 From the **Load Balancing** drop-down menu, specify how the standard or distributed switch selects an uplink to handle traffic from a virtual machine or VMkernel adapter.

Option	Description
Route based on the originating virtual port	Select an uplink based on the virtual port where the traffic entered the virtual switch.
Route based on IP hash	Select an uplink based on a hash of the source and destination IP addresses of each packet. For non-IP packets, the switch simply uses the data at those fields to compute the hash. IP-based teaming requires that the physical switch be configured with EtherChannel.
Route based on source MAC hash	Select an uplink based on a hash of the source Ethernet.
Route based on physical NIC load	Available for distributed port groups or distributed ports. Select an uplink based on the current load of the physical network adapters connected to the port group or port. If an uplink remains busy at 75% or higher for 30 seconds, the host proxy switch moves a part of the virtual machine traffic to a physical adapter that has free capacity.
Use explicit failover order	From the list of active adapters, always use the highest order uplink that passes failover detection criteria.

7 From the **Network Failover Detection** drop-down menu, specify the method that the standard or distributed switch uses for failover detection.

Option	Description
Link Status only	Relies only on the link status that the network adapter provides. This option detects failures, such as removed cables and physical switch power failures. However, it does not detect configuration errors, such as the following: ■ Physical switch port that is blocked by spanning tree or is misconfigured to the wrong VLAN. ■ Pulled cable that connects a physical switch to another networking devices, for example, an upstream switch.
Beacon Probing	Sends out and listens for beacon probes on all NICs in the team and uses this information, in addition to link status, to determine link failure. ESX/ESXi sends beacon packets every second. Beaconing is most useful with three or more NICs in a team because ESX/ESXi can detect failures of a single adapter. If only two NICs are assigned and one of them loses connectivity, the switch cannot determine which NIC needs to be taken out of service because both do not receive beacons and as a result all packets sent to both uplinks. Using at least three NICs in such a team allows for n–2 failures where *n* is the number of NICs in the team before reaching an ambiguous situation. The NICs must be in an active/active or active/standby configuration because the NICs in an unused state do not participate in beacon probing.

8 From the **Notify Switches** drop-down menu, select whether the standard or distributed switch notifies the physical switch in case of a failover.

If you select **Yes**, whenever a virtual NIC is connected to the virtual 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 physical switches. Notifying the physical switch offers lowest latency when a failover or a migration with vSphere vMotion occurs.

NOTE Set this option to **No** if a connected virtual machines is using Microsoft Network Load Balancing in unicast mode. No issues exist with Network Load Balancing running in multicast mode.

9 From the **Failback** drop-down menu, determine whether a physical adapter is returned to active status after recovering from a failure.

If failback is set to **Yes** (default), the adapter is returned to active duty immediately upon recovery, displacing the standby adapter that took over its slot, if any.

If failback is set to **No** to a standard port, a failed adapter is left inactive after recovery until another currently active adapter fails and must be replaced.

10 Specify how the uplinks in a team are used when a failover occurs by configuring the Failover Order list.

If you want to use some uplinks but reserve others for emergencies in case the uplinks in use fail, use the up and down arrow keys to move them into different groups.

Option	Description
Active adapters	Continue to use the uplink if the network adapter connectivity is up and active.
Standby adapters	Use this uplink if one of the active physical adapter is down.
Unused adapters	Do not use this uplink.

11 Click **OK**.

Edit the Teaming and Failover Policy on a Distributed Port Group in the vSphere Web Client

The teaming and failover policy on a distributed port group lets you determine how network traffic of a group of virtual machines or VMkernel adapters is distributed between the uplink adapters in a vSphere Distributed Switch and how to reroute traffic in the event of an adapter failure.

Procedure

1 In the vSphere Web Client, navigate to the distributed switch.

2 Right-click the distributed switch, and select **Manage Distributed Port Groups**.

3 Select the **Teaming and failover** check box and click **Next**.

4 Select the port group that you want to configure and click **Next**.

5 From the **Load Balancing** drop-down menu, specify how the standard or distributed switch selects an uplink to handle traffic from a virtual machine or VMkernel adapter.

Option	Description
Route based on the originating virtual port	Select an uplink based on the virtual port where the traffic entered the virtual switch.
Route based on IP hash	Select an uplink based on a hash of the source and destination IP addresses of each packet. For non-IP packets, the switch simply uses the data at those fields to compute the hash. IP-based teaming requires that the physical switch be configured with EtherChannel.
Route based on source MAC hash	Select an uplink based on a hash of the source Ethernet.
Route based on physical NIC load	Available for distributed port groups or distributed ports. Select an uplink based on the current load of the physical network adapters connected to the port group or port. If an uplink remains busy at 75% or higher for 30 seconds, the host proxy switch moves a part of the virtual machine traffic to a physical adapter that has free capacity.
Use explicit failover order	From the list of active adapters, always use the highest order uplink that passes failover detection criteria.

6 From the **Network Failover Detection** drop-down menu, specify the method that the standard or distributed switch uses for failover detection.

Option	Description
Link Status only	Relies only on the link status that the network adapter provides. This option detects failures, such as removed cables and physical switch power failures. However, it does not detect configuration errors, such as the following: ■ Physical switch port that is blocked by spanning tree or is misconfigured to the wrong VLAN. ■ Pulled cable that connects a physical switch to another networking devices, for example, an upstream switch.
Beacon Probing	Sends out and listens for beacon probes on all NICs in the team and uses this information, in addition to link status, to determine link failure. ESX/ESXi sends beacon packets every second. Beaconing is most useful with three or more NICs in a team because ESX/ESXi can detect failures of a single adapter. If only two NICs are assigned and one of them loses connectivity, the switch cannot determine which NIC needs to be taken out of service because both do not receive beacons and as a result all packets sent to both uplinks. Using at least three NICs in such a team allows for n–2 failures where *n* is the number of NICs in the team before reaching an ambiguous situation. The NICs must be in an active/active or active/standby configuration because the NICs in an unused state do not participate in beacon probing.

7 From the **Notify Switches** drop-down menu, select whether the standard or distributed switch notifies the physical switch in case of a failover.

If you select **Yes**, whenever a virtual NIC is connected to the virtual 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 physical switches. Notifying the physical switch offers lowest latency when a failover or a migration with vSphere vMotion occurs.

NOTE Set this option to **No** if a connected virtual machines is using Microsoft Network Load Balancing in unicast mode. No issues exist with Network Load Balancing running in multicast mode.

8 From the **Failback** drop-down menu, determine whether a physical adapter is returned to active status after recovering from a failure.

If failback is set to **Yes** (default), the adapter is returned to active duty immediately upon recovery, displacing the standby adapter that took over its slot, if any.

If failback is set to **No** to a distributed port, a failed adapter is left inactive after recovery only if the associated virtual machine is running. If the **Failback** option is **No** and a virtual machine is powered off when all active physical adapters fail and then one of them recovers, after the virtual machine is powered on, the virtual NIC is connected to the recovered adapter instead of to a standby one. Powering a virtual machine off and then on leads to reconnecting the virtual NIC to a distributed port. The distributed switch considers the port as newly added, and assigns it the default uplink port, that is, the active uplink adapter.

9 Specify how the uplinks in a team are used when a failover occurs by configuring the Failover Order list.

If you want to use some uplinks but reserve others for emergencies in case the uplinks in use fail, use the up and down arrow keys to move them into different groups.

Option	Description
Active adapters	Continue to use the uplink if the network adapter connectivity is up and active.
Standby adapters	Use this uplink if one of the active physical adapter is down.
Unused adapters	Do not use this uplink.

10 Review your settings and click **Finish**.

Edit Distributed Port Teaming and Failover Policies with the vSphere Web Client

The teaming and failover policy lets you determine how network traffic is distributed between physical adapters and how traffic is rerouted in the event of an adapter failure.

Prerequisites

Enable the port-level override option for this policy. See "Edit Advanced Distributed Port Group Settings with the vSphere Web Client," on page 44.

Procedure

1 In the vSphere Web Client, navigate to the distributed switch.

2 On the **Manage** tab, and select **Ports**.

3 Select a distributed port from the list and click **Edit distributed port settings**.

4 Click **Teaming and failover** and select the check box next to the properties that you want to override.

5 From the **Load Balancing** drop-down menu, specify how the standard or distributed switch selects an uplink to handle traffic from a virtual machine or VMkernel adapter.

Option	Description
Route based on the originating virtual port	Select an uplink based on the virtual port where the traffic entered the virtual switch.
Route based on IP hash	Select an uplink based on a hash of the source and destination IP addresses of each packet. For non-IP packets, the switch simply uses the data at those fields to compute the hash. IP-based teaming requires that the physical switch be configured with EtherChannel.
Route based on source MAC hash	Select an uplink based on a hash of the source Ethernet.
Route based on physical NIC load	Available for distributed port groups or distributed ports. Select an uplink based on the current load of the physical network adapters connected to the port group or port. If an uplink remains busy at 75% or higher for 30 seconds, the host proxy switch moves a part of the virtual machine traffic to a physical adapter that has free capacity.
Use explicit failover order	From the list of active adapters, always use the highest order uplink that passes failover detection criteria.

6 From the **Network Failover Detection** drop-down menu, specify the method that the standard or distributed switch uses for failover detection.

Option	Description
Link Status only	Relies only on the link status that the network adapter provides. This option detects failures, such as removed cables and physical switch power failures. However, it does not detect configuration errors, such as the following: ■ Physical switch port that is blocked by spanning tree or is misconfigured to the wrong VLAN. ■ Pulled cable that connects a physical switch to another networking devices, for example, an upstream switch.
Beacon Probing	Sends out and listens for beacon probes on all NICs in the team and uses this information, in addition to link status, to determine link failure. ESX/ESXi sends beacon packets every second. Beaconing is most useful with three or more NICs in a team because ESX/ESXi can detect failures of a single adapter. If only two NICs are assigned and one of them loses connectivity, the switch cannot determine which NIC needs to be taken out of service because both do not receive beacons and as a result all packets sent to both uplinks. Using at least three NICs in such a team allows for n–2 failures where *n* is the number of NICs in the team before reaching an ambiguous situation. The NICs must be in an active/active or active/standby configuration because the NICs in an unused state do not participate in beacon probing.

7 From the **Notify Switches** drop-down menu, select whether the standard or distributed switch notifies the physical switch in case of a failover.

If you select **Yes**, whenever a virtual NIC is connected to the virtual 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 physical switches. Notifying the physical switch offers lowest latency when a failover or a migration with vSphere vMotion occurs.

NOTE Set this option to **No** if a connected virtual machines is using Microsoft Network Load Balancing in unicast mode. No issues exist with Network Load Balancing running in multicast mode.

8 From the **Failback** drop-down menu, determine whether a physical adapter is returned to active status after recovering from a failure.

If failback is set to **Yes** (default), the adapter is returned to active duty immediately upon recovery, displacing the standby adapter that took over its slot, if any.

If failback is set to **No** to a distributed port, a failed adapter is left inactive after recovery only if the associated virtual machine is running. If the **Failback** option is **No** and a virtual machine is powered off when all active physical adapters fail and then one of them recovers, after the virtual machine is powered on, the virtual NIC is connected to the recovered adapter instead of to a standby one. Powering a virtual machine off and then on leads to reconnecting the virtual NIC to a distributed port. The distributed switch considers the port as newly added, and assigns it the default uplink port, that is, the active uplink adapter.

9 Specify how the uplinks in a team are used when a failover occurs by configuring the Failover Order list.

If you want to use some uplinks but reserve others for emergencies in case the uplinks in use fail, use the up and down arrow keys to move them into different groups.

Option	Description
Active adapters	Continue to use the uplink if the network adapter connectivity is up and active.
Standby adapters	Use this uplink if one of the active physical adapter is down.
Unused adapters	Do not use this uplink.

10 Click **OK**.

VLAN Policy

VLAN policies determine how VLANs function across your network environment.

A virtual local area network (VLAN) is a group of hosts with a common set of requirements, which communicate as if they were attached to the same broadcast domain, regardless of their physical location. A VLAN has the same attributes as a physical local area network (LAN), but it allows for end stations to be grouped together even if not on the same network switch.

The scope of VLAN policies can be distributed port groups and ports, and uplink port groups and ports.

Edit the VLAN Policy on a Distributed Port Group in the vSphere Web Client

Set the VLAN policy on a distributed port group to apply VLAN tagging globally on all distributed ports .

Procedure

1 Browse to a distributed switch in the vSphere Web Client navigator.

2 Right-click the distributed switch in the navigator and select **Manage Distributed Port Groups**.

3 Select the **VLAN** check box and click **Next**.

4 Select the port group that you want to edit and click **Next**.

5 Select the type of VLAN traffic filtering and marking from the **Type** drop-down menu and click **Next**.

Option	Description
None	Do not use VLAN. Use this option in case of External Switch Tagging (EST).
VLAN	Tag traffic with the ID from the **VLAN ID** field. Type a number between 1 and 4094 for Virtual Switch Tagging (VST).
VLAN Trunking	Pass VLAN traffic with ID within the **VLAN trunk range**. You can set multiple ranges and individual VLANs by using a comma-separated list. Use this option in VGT.
Private VLAN	Associate the traffic with a private VLAN created on the distributed switch.

6 Review your settings and click **Finish**.

Edit the VLAN Policy on a Distributed Port with the vSphere Web Client

Use the VLAN policy on a distributed port to integrate the virtual traffic through the port with physical VLANs in a different way from the parent distributed port group.

Prerequisites

To override the VLAN policy at the port level, enable the port-level overrides. See "Edit Advanced Distributed Port Group Settings with the vSphere Web Client," on page 44

Procedure

1 Browse to a distributed switch in the vSphere Web Client navigator.

2 Click the **Manage** tab and select **Ports**.

3 Select a port from the list.

4 Click **Edit distributed port settings**.

5 Click **VLAN** and select **Override**.

Configure VLAN traffic through the distributed port from the **VLAN Type** drop-down meny.

Option	Description
None	Do not use VLAN. Use this option in case of External Switch Tagging (EST).
VLAN	Tag traffic with the ID from the **VLAN ID** field. Type a number between 1 and 4094 for Virtual Switch Tagging (VST).
VLAN Trunking	Pass VLAN traffic with ID within the **VLAN trunk range**. You can set multiple ranges and individual VLANs by using a comma-separated list. Use this option in VGT.
Private VLAN	Associate the traffic with a private VLAN created on the distributed switch.

6 Click **OK**.

Edit the VLAN Policy on an Uplink Port Group in the vSphere Web Client

Set the VLAN policy on an uplink port group to configure VLAN traffic processing generally for all member uplinks.

Use the VLAN policy at the uplink port level to propagate a trunk range of VLAN IDs to the physical network adapters for traffic filtering. The physical network adapters drop the packets from the other VLANs if the adapters support filtering by VLAN. Setting a trunk range improves networking performance because physical network adapters filter traffic instead of the uplink ports in the group.

If you have a physical network adapter which does not support VLAN filtering, the VLANs might still not be blocked. In this case, configure VLAN filtering on a distributed port group or on a distributed port.

See the technical documentation from the adapter vendors for information about VLAN filtering support.

Procedure

1 Locate an uplink port group in the vSphere Web Client.

a To locate an uplink port group, select a distributed switch and click the **Related Objects** tab.

b Select the **Uplink Port Groups** tab and locate the uplink group in the list.

2 Right-click the uplink port group in the list and select **Edit Settings**.

3 Click **VLAN** and type a **VLAN trunk range** to propagate to the physical network adapters.

For trunking of several ranges and individual VLANs, separate the entries with commas.

4 Click **OK**.

Edit the VLAN Policy on an Uplink Port with the vSphere Web Client

Set the VLAN policy on an uplink port to handle VLAN traffic through the port in a different way than for the parent uplink port group.

Use the VLAN policy at the uplink port to propagate a trunk range of VLAN IDs to the physical network adapter for traffic filtering. The physical network adapter drops packets from the other VLANs if the adapter supports filtering by VLAN. Setting a trunk range improves networking performance because the physical network adapter filters traffic instead of the uplink port.

If you have a physical network adapter which does not support VLAN filtering, the VLANs might still not be blocked. In this case, configure VLAN filtering on a distributed port group or on a distributed port.

See the technical documentation from the adapter vendor for information about VLAN filtering support.

Prerequisites

To override the VLAN policy at the port level, enable the port-level overrides. See "Edit Advanced Distributed Port Group Settings with the vSphere Web Client," on page 44.

Procedure

1 Locate an uplink port group in the vSphere Web Client.

a To locate an uplink port group, select a distributed switch and click the **Related Objects** tab.

b Select the **Uplink Port Groups** tab and double-click an uplink port group from the list.

The uplink port group appears at the top level of the navigator on the left.

2 Click the **Manage** tab and select **Ports**.

3 Select an uplink port from the list and click **Edit distributed port settings**.

4 Click **VLAN** and select the **Override** check box.

5 Type a **VLAN trunk range** to propagate to the physical network adapter.

For trunking of several ranges and individual VLANs, separate the entries with commas.

6 Click **OK**.

Security Policy

Networking security policy provides protection of traffic against MAC address impersonation and unwanted port scanning

The security policy of a standard or distributed switch is implemented in Layer 2 (Data Link Layer) of the network protocol stack. The three elements of the security policy are promiscuous mode, MAC address changes, and forged transmits. See the *vSphere Security* documentation for information about potential networking threats.

Edit Security Policy for a vSphere Standard Switch in the vSphere Web Client

For a vSphere standard switch you can configure security policy to reject MAC address and promiscuous mode changes in the guest operating system of a virtual machine.

You can configure security policy for an individual standard port group as well.

Procedure

1 In the vSphere Web Client, navigate to the host.

2 On the **Manage** tab, click **Networking**, and select **Virtual switches**.

3 Select a standard switch from the list and click **Edit settings**.

4 Select **Security**, and reject or accept promiscuous mode activation or MAC address changes in the guest operating system of the virtual machines attached to the standard switch.

By default, enabling promiscuous mode and MAC address changes for both inbound and outbound traffic is not accepted.

Option	Description
Promiscuous mode	■ **Reject**: Placing an adapter in promiscuous mode from the guest operating system does not result in receiving frames for other virtual machines. ■ **Accept**: If an adapter is placed in promiscuous mode from the guest operating system, the switch allows the guest adapter to receive all frames passed on the switch in compliance with the active VLAN policy for the port to which the adapter is connected. Firewalls, port scanners, intrusion detection systems and so on, need to run in promiscuous mode.
MAC address changes	■ **Reject**: If you set the **MAC Address Changes** to **Reject** and the guest operating system changes the MAC address of the adapter to a value different from the address in the virtual machine configuration file (`.vmx`), the switch drops all inbound frames to the virtual machine adapter. If the guest operating system changes the MAC address back, the virtual machine receives frames again. ■ **Accept**: If the guest operating system changes the MAC address of a network adapter, the switch allows frames to the new address of the adapter to pass.
Forged transmits	■ **Reject**: The switch drops any outbound frame from a virtual machine adapter with a source MAC address that is different from the one in the `.vmx` configuration file. ■ **Accept**: The switch does not perform filtering and permits all outbound frames.

5 Click **OK**.

Edit the Layer 2 Security Policy Exception for a Standard Port Group in the vSphere Web Client

By using the security policy on a port group you can accept or reject promiscuous mode and MAC address changes in the guest operating system of a virtual machine connected to the group.

Procedure

1 In the vSphere Web Client, navigate to the host.

2 On the **Manage** tab, click **Networking**, and select **Virtual switches**.

3 Select a standard switch from the list.

The topology diagram of the standard switch appears.

4 In the topology diagram of the standard switch, click the name of the standard port group to configure.

5 Click **Edit settings**.

6 In the **Security** section, select the check boxes next to the security policies to override and use the drop-down menus to configure security over the virtual machine traffic through the ports of the group.

By default, enabling promiscuous mode and MAC address changes for both inbound and outbound traffic is not accepted.

Option	Description
Promiscuous mode	■ **Reject**: Placing an adapter in promiscuous mode from the guest operating system does not result in receiving frames for other virtual machines. ■ **Accept**: If an adapter is placed in promiscuous mode from the guest operating system, the switch allows the guest adapter to receive all frames passed on the switch in compliance with the active VLAN policy for the port to which the adapter is connected. Firewalls, port scanners, intrusion detection systems and so on, need to run in promiscuous mode.
MAC address changes	■ **Reject**: If you set the **MAC Address Changes** to **Reject** and the guest operating system changes the MAC address of the adapter to a value different from the address in the virtual machine configuration file (`.vmx`), the switch drops all inbound frames to the virtual machine adapter. If the guest operating system changes the MAC address back, the virtual machine receives frames again. ■ **Accept**: If the guest operating system changes the MAC address of a network adapter, the switch allows frames to the new address of the adapter to pass.
Forged transmits	■ **Reject**: The switch drops any outbound frame from a virtual machine adapter with a source MAC address that is different from the one in the `.vmx` configuration file. ■ **Accept**: The switch does not perform filtering and permits all outbound frames.

7 Click **OK**.

Edit the Security Policy for a Distributed Port Group in the vSphere Web Client

You can set a security policy on a distributed port group to allow or reject enabling promiscuous mode and MAC address changes from the guest operating system of a virtual machine associated with the port group.

You can configure security policy for an individual distributed port as well.

Procedure

1 Navigate to a distributed switch in the vSphere Web Client .

2 Right-click the distributed switch and select **Manage Distributed Port Groups**.

3 Click the **Security** check box and click **Next**.

4 Select the distributed port group to configure and click **Next**.

5 Use the drop-down menus to edit the security settings on the traffic through the ports of the group, and click **Next**.

By default, enabling promiscuous mode and MAC address changes for both inbound and outbound traffic is not accepted.

Option	Description
Promiscuous mode	■ **Reject**: Placing an adapter in promiscuous mode from the guest operating system does not result in receiving frames for other virtual machines. ■ **Accept**: If an adapter is placed in promiscuous mode from the guest operating system, the switch allows the guest adapter to receive all frames passed on the switch in compliance with the active VLAN policy for the port to which the adapter is connected. Firewalls, port scanners, intrusion detection systems and so on, need to run in promiscuous mode.
MAC address changes	■ **Reject**: If you set the **MAC Address Changes** to **Reject** and the guest operating system changes the MAC address of the adapter to a value different from the address in the virtual machine configuration file (`.vmx`), the switch drops all inbound frames to the virtual machine adapter. If the guest operating system changes the MAC address back, the virtual machine receives frames again. ■ **Accept**: If the guest operating system changes the MAC address of a network adapter, the switch allows frames to the new address of the adapter to pass.
Forged transmits	■ **Reject**: The switch drops any outbound frame from a virtual machine adapter with a source MAC address that is different from the one in the `.vmx` configuration file. ■ **Accept**: The switch does not perform filtering and permits all outbound frames.

6 Review your settings and click **Finish**.

Edit Distributed Port Security Policies with the vSphere Web Client

On a distributed port you can override the policy inherited from the distributed port group to accept or reject enabling promiscuous mode and MAC address changes in the guest operating system of the virtual machine connected to the port.

Prerequisites

Enable port-level overrides. See "Edit Advanced Distributed Port Group Settings with the vSphere Web Client," on page 44

Procedure

1 Navigate to a distributed switch in the vSphere Web Client.

2 Click the **Manage** tab, and select **Ports**.

3 Select a port from the list.

4 Click **Edit distributed port settings**.

5 Click **Security** and select the check box for the settings you want to override for the traffic through the port.

By default, enabling promiscuous mode and MAC address changes for both inbound and outbound traffic is not accepted.

Option	Description
Promiscuous mode	■ **Reject**: Placing an adapter in promiscuous mode from the guest operating system does not result in receiving frames for other virtual machines. ■ **Accept**: If an adapter is placed in promiscuous mode from the guest operating system, the switch allows the guest adapter to receive all frames passed on the switch in compliance with the active VLAN policy for the port to which the adapter is connected. Firewalls, port scanners, intrusion detection systems and so on, need to run in promiscuous mode.
MAC address changes	■ **Reject**: If you set the **MAC Address Changes** to **Reject** and the guest operating system changes the MAC address of the adapter to a value different from the address in the virtual machine configuration file (`.vmx`), the switch drops all inbound frames to the virtual machine adapter. If the guest operating system changes the MAC address back, the virtual machine receives frames again. ■ **Accept**: If the guest operating system changes the MAC address of a network adapter, the switch allows frames to the new address of the adapter to pass.
Forged transmits	■ **Reject**: The switch drops any outbound frame from a virtual machine adapter with a source MAC address that is different from the one in the `.vmx` configuration file. ■ **Accept**: The switch does not perform filtering and permits all outbound frames.

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.

Average Bandwidth Establishes the number of bits per second to allow across a port, averaged over time. This number is the allowed average load.

Peak Bandwidth 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.

Burst Size 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 in the vSphere Web Client

ESXi lets you to shape outbound traffic on standard switches. The traffic shaper restricts the network bandwidth available to any port, but you can also configure it 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.

Average Bandwidth — Establishes the number of bits per second to allow across a port, averaged over time (the allowed average load).

Peak Bandwidth — The maximum number of bits per second to allow across a port when it is sending 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.

Burst Size — The maximum number of bytes to allow in a burst. If this parameter is set, a port might gain a burst bonus when it does not use all its allocated bandwidth. Whenever the port needs more bandwidth than specified by **Average Bandwidth**, it might be allowed to temporarily transmit data at a higher speed if a burst bonus is available. This parameter tops the number of bytes that can be accumulated in the burst bonus and transferred at a higher speed.

Procedure

1 In the vSphere Web Client, navigate to the host.

2 On the **Manage** tab, click **Networking**, and select **Virtual switches**.

3 Select a standard switch from the list and click **Edit settings**.

4 Click **Traffic shaping** and enable or disable traffic shaping policy exceptions with the **Status** drop-down menu.

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

5 For each traffic shaping policy (**Average Bandwidth**, **Peak Bandwidth**, and **Burst Size**), enter a bandwidth value.

6 Click **OK**.

Edit the Traffic Shaping Policy for a Standard Port Group in the vSphere Web Client

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

Prerequisites

Enable the port-level overrides. See "Edit Advanced Distributed Port Group Settings with the vSphere Web Client," on page 44

Procedure

1 Browse to a host in the vSphere Web Client navigator.

2 Click the **Manage** tab, and click **Networking > Virtual Switches**.

3 Select a standard switch from the list.

A schematic of the standard switch infrastructure appears.

4 Click **Edit settings**.

5 Click **Traffic Shaping** and click the **Override** check box to override the traffic shaping policy at the standard port group level and enter settings.

NOTE If you have not enabled port group-level overrides, the options are not available.

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

6 Click **OK**.

Edit the Traffic Shaping Policy on a Distributed Port Group in the vSphere Web Client

ESXi allows you to shape both inbound and outbound traffic on vSphere distributed port groups. The traffic shaper restricts the network bandwidth available to any port in the group, but might 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 Browse to a distributed switch in the vSphere Web Client navigator.

2 Right-click the distributed switch in the navigator and select **Manage Distributed Port Groups**.

3 Select the **Traffic Shaping** check box and click **Next**.

4 On the **Select port groups** page, select a port group from the list and click **Next**

5 Click **Traffic shaping**, and select the **Override** check box to override ingress traffic shaping, egress traffic shaping, or both.

NOTE The traffic is classified to ingress and egress according to the traffic direction in the switch, not in the host.

Option	Description
Status	If you enable either **Ingress Traffic Shaping** or **Egress Traffic Shaping** by using the **Status** drop-down menus, you are setting limits on the amount of networking bandwidth allocated for each VMkernel adapter or virtual network adapter associated with this particular port group. If you disable the policy, services have a free, clear connection to the physical network by default.
Average Bandwidth	Establishes the number of bits per second to allow across a port, averaged over time, that is, the allowed average load.

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

6 Review your settings and click **Finish**.

Use the **Back** button to edit any settings.

Edit the Traffic Shaping Policy on a Distributed Port in the vSphere Web Client

ESXi allows you to shape both inbound and outbound traffic on vSphere distributed switches. The traffic shaper restricts the network bandwidth available to a port, but might also be configured to temporarily allow bursts of traffic to flow through the port at higher speeds.

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

Prerequisites

Enable the port-level override option for this policy. See "Edit Advanced Distributed Port Group Settings with the vSphere Web Client," on page 44.

Procedure

1 Browse to a distributed switch in the vSphere Web Client.

2 To navigate to the distributed ports of the distributed switch, click **Manage > Ports**.

3 Select a port from the list.

4 Click **Edit distributed port settings**.

5 Click **Traffic shaping**, and select the **Override** check box to override ingress traffic shaping, egress traffic shaping, or both.

NOTE The traffic is classified to ingress and egress according to the traffic direction in the switch, not in the host.

Option	Description
Status	If you enable either **Ingress Traffic Shaping** or **Egress Traffic Shaping** by using the **Status** drop-down menus, you are setting limits on the amount of networking bandwidth allocated for each VMkernel adapter or virtual network adapter associated with this particular port group. If you disable the policy, services have a free, clear connection to the physical network by default.
Average Bandwidth	Establishes the number of bits per second to allow across a port, averaged over time, that is, the allowed average load.

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

6 Review your settings in the **Ready to complete** section and click **Finish**.

Use the **Back** button to edit any settings.

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 121.

Edit the Resource Allocation Policy on a Distributed Port Group in the vSphere Web Client

Associate a distributed port group with a network resource pool to give you greater control over the bandwidth that is 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 Browse to a distributed switch in the vSphere Web Client navigator.

2 Right-click the distributed switch in the navigator and select **Manage Distributed Port Groups**.

3 Select the **Resource allocation** check box and click **Next**.

4 Select the distributed port group to edit and click **Next**.

5 Add or remove the distributed port group from the network resource pool and click **Next**.

- To add the distributed port group, select a user-defined resource pool from the **Network resource pool** drop-down menu.
- To remove the distributed port group, select **default** from the **Network resource pool** drop-down menu.

6 Review your settings in the **Ready to complete** section and click **Finish**.

Use the **Back** button to change any settings.

Edit the Resource Allocation Policy on a Distributed Port in the vSphere Web Client

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.
- Enable port-level overrides. See "Edit Advanced Distributed Port Group Settings with the vSphere Web Client," on page 44.

Procedure

1. Browse to a distributed switch in the vSphere Web Client navigator.
2. Click the **Manage** tab and click **Ports**.
3. Select a port from the list and click **Edit distributed port settings**.
4. In the **Properties** section, click the **Override** check box and add or remove the port from a network resource pool.

   If you did not enable port-level overrides, the options are not available.

   - To **add** the distributed port to a resource pool, select a user-defined resource pool from the **Network resource pool** drop-down menu.
   - To **remove** the distributed port from a resource pool, select **Default** from the **Network resource pool** drop-down menu.
5. 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 with the vSphere Web Client," on page 159.

Edit the Monitoring Policy on a Distributed Port Group in the vSphere Web Client

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

Procedure

1. Browse to a distributed switch in the vSphere Web Client navigator.
2. Right-click the distributed switch in the object navigator and select **Manage Distributed Port Groups**.
3. Select the **Monitoring** check box and click **Next**.
4. Select the distributed port group to edit and click **Next**.

5 Use the drop-menu to enable or disable NetFlow and click **Next**.

Option	Description
Disabled	NetFlow is disabled on the distributed port group.
Enabled	NetFlow is enabled on the distributed port group. You can configure NetFlow settings at the vSphere distributed switch level. See "Configure NetFlow Settings with the vSphere Web Client," on page 159.

6 Review your settings and click **Finish**.

Use the **Back** button to change any settings.

Edit the Monitoring Policy on a Distributed Port in the vSphere Web Client

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

Prerequisites

To override the monitoring policy at the port level, enable the port-level overrides. See "Edit Advanced Distributed Port Group Settings with the vSphere Web Client," on page 44

Procedure

1 Browse to a distributed switch in the vSphere Web Client navigator.

2 Click the **Manage** tab, and click **Ports**.

3 Select a port from the list.

Detailed port setting information appears at the bottom of the screen.

4 Click **Edit distributed port settings**.

5 Click **Monitoring** and click the check box to override the NetFlow settings at the port group level.

6 Enable or disable Netflow from the drop-down menu.

NOTE If you have no enabled port-level overrides, the options are not available.

Option	Description
Disabled	NetFlow is disabled on the distributed port group.
Enabled	NetFlow is enabled on the distributed port group. You can configure NetFlow settings at the vSphere distributed switch level. See "Configure NetFlow Settings with the vSphere Web Client," on page 159.

7 Click **OK**.

Traffic Filtering and Marking Policy

In a vSphere distributed switch 5.5 and later, by using the traffic filtering and marking policy, you can protect the virtual network from unwanted traffic and security attacks or apply a QoS tag to a certain type of traffic.

The traffic filtering and marking policy represents an ordered set of network traffic rules for security and for QoS tagging of the data flow through the ports of a distributed switch. In general, a rule consists of a qualifier for traffic, and of an action for restricting or prioritizing the matching traffic.

The vSphere distributed switch applies rules on traffic at different places in the data stream. The distributed switch applies traffic filter rules on the data path between the virtual machine network adapter and distributed port, or between the uplink port and physical network adapter for rules on uplinks.

Traffic Filtering and Marking on a Distributed Port Group or Uplink Port Group in the vSphere Web Client

Set traffic rules at the level of distributed port groups or of uplink port groups to introduce filtering and priority tagging for traffic access over virtual machines, VMkernel adapters, or physical adapters.

- Enable Traffic Filtering and Marking on a Distributed Port Group or Uplink Port Group in the vSphere Web Client on page 99

  Enable the traffic filtering and marking policy on a port group if you want to configure traffic security and marking on all virtual machine network adapters or uplink adapters that are participating in the group.

- Mark Traffic on a Distributed Port Group or Uplink Port Group in the vSphere Web Client on page 100

  Assign priority tags to traffic, such as VoIP and streaming video, that has higher networking requirements for bandwidth, low latency, and so on. You can mark the traffic with a CoS tag in Layer 2 of the network protocol stack or with a DSCP tag in Layer 3.

- Filter Traffic on a Distributed Port Group or Uplink Port Group in the vSphere Web Client on page 102

  Allow or stop traffic for securing the data that flows through the ports of a distributed port group or uplink port group.

- Working with Network Traffic Rules on a Distributed Port Group or Uplink Port Group in the vSphere Web Client on page 103

  Define traffic rules in a distributed port group or uplink port group to introduce a policy for processing traffic related to virtual machines or to physical adapters. You can filter specific traffic or describe its QoS demands.

- Disable Traffic Filtering and Marking on a Distributed Port Group or Uplink Port Group in the vSphere Web Client on page 105

  Let traffic flow to virtual machines or physical adapters without additional control related to security or QoS by disabling the traffic filtering and marking policy.

Enable Traffic Filtering and Marking on a Distributed Port Group or Uplink Port Group in the vSphere Web Client

Enable the traffic filtering and marking policy on a port group if you want to configure traffic security and marking on all virtual machine network adapters or uplink adapters that are participating in the group.

NOTE You can disable the traffic filtering and marking policy on a particular port to avoid processing the traffic flowing through the port. See "Disable Traffic Filtering and Marking on a Distributed Port or Uplink Port in the vSphere Web Client," on page 111.

Procedure

1. Locate a distributed port group or an uplink port group in the vSphere Web Client.
   a. Select a distributed switch and click the **Related Objects** tab.
   b. Click **Distributed Port Groups** to see the list of distributed port groups, or click **Uplink Port Groups** to see the list of uplink port groups.
2. Right-click the port group and select **Edit settings**.
3. Select **Traffic filtering and marking**.
4. From the **Status** drop-down menu, select **Enabled**.

5 Click **OK**.

What to do next

Set up traffic marking or filtering on the data that is flowing through the ports of the distributed port group or through the uplink port group. See "Mark Traffic on a Distributed Port Group or Uplink Port Group in the vSphere Web Client," on page 100 and "Filter Traffic on a Distributed Port Group or Uplink Port Group in the vSphere Web Client," on page 102.

Mark Traffic on a Distributed Port Group or Uplink Port Group in the vSphere Web Client

Assign priority tags to traffic, such as VoIP and streaming video, that has higher networking requirements for bandwidth, low latency, and so on. You can mark the traffic with a CoS tag in Layer 2 of the network protocol stack or with a DSCP tag in Layer 3.

Priority tagging is a mechanism to mark traffic that has higher QoS demands. In this way, the network can recognize different classes of traffic. The network devices can handle the traffic from each class according to its priority and requirements.

You can also re-tag traffic to either raise or lower the importance of the flow. By using a low QoS tag, you can restrict data tagged in a guest operating system.

Procedure

1 Locate a distributed port group or an uplink port group in the vSphere Web Client.

   a Select a distributed switch and click the **Related Objects** tab.

   b Click **Distributed Port Groups** to see the list of distributed port groups, or click **Uplink Port Groups** to see the list of uplink port groups.

2 Right-click the port group and select **Edit settings**.

3 Select **Traffic filtering and marking**.

4 If traffic filtering and marking is disabled, enable it from the **Status** drop-down menu.

5 Click **New** to create a new rule, or select a rule and click **Edit** to edit it.

6 In the network traffic rule dialog box, select the **Tag** option from the **Action** drop-down menu.

7 Set the priority tag for the traffic within the scope of the rule.

Option	Description
CoS value	Mark the traffic matching the rule with a CoS priority tag in network Layer 2. Select **Update CoS tag** and type a value from 0 to 7.
DSCP value	Mark the traffic associated with the rule with a DSCP tag in network Layer 3. Select **Update DSCP value** and type a value from 0 to 63.

8 Specify the kind of traffic that the rule is applicable to.

To determine if a data flow is in the scope of a rule for marking or filtering, the vSphere distributed switch examines the direction of the traffic, and properties like source and destination, VLAN, next level protocol, infrastructure traffic type, and so on.

a From the **Traffic direction** drop-down menu, select whether the traffic must be ingress, egress, or both so that the rule recognizes it as matching.

The direction also influences how you are going to identify the traffic source and destination.

b By using qualifiers for system data type, Layer 2 packet attributes, and Layer 3 packet attributes set the properties that packets must have to match the rule.

A qualifier represents a set of matching criteria related to a networking layer. You can match traffic to system data type, Layer 2 traffic properties, and Layer 3 traffic properties. You can use the qualifier for a specific networking layer or can combine qualifiers to match packets more precisely.

- Use the system traffic qualifier to match packets to the type of virtual infrastructure data that is flowing through the ports of the group . For example, you can select NFS for data transfers to network storage.
- Use the MAC traffic qualifier to match packets by MAC address, VLAN ID, and next level protocol.

  Locating traffic with a VLAN ID on a distributed port group works with Virtual Guest Tagging (VGT). To match traffic to VLAN ID if Virtual Switch Tagging (VST) is active, use a rule on an uplink port group or uplink port.
- Use the IP traffic qualifier to match packets by IP version, IP address, and next level protocol and port.

9 In the rule dialog box, click **OK** to save the rule.

Example: Voice over IP Traffic Marking

Voice over IP (VoIP) flows have special requirements for QoS in terms of low loss and delay. The traffic related to the Session Initiation Protocol (SIP) for VoIP usually has a DSCP tag equal to 26, which stands for Assured Forwarding Class 3 with Low Drop Probability (AF31).

For example, to mark outgoing SIP UDP packets to a subnet 192.168.2.0/24, you can use the following rule:

Rule Parameter	Parameter Value
Action	Tag
DSCP value	26
Traffic direction	Egress
Traffic qualifiers	IP Qualifier
Protocol	UDP
Destination port	5060
Source address	IP address matches 192.168.2.0 with prefix length 24

Filter Traffic on a Distributed Port Group or Uplink Port Group in the vSphere Web Client

Allow or stop traffic for securing the data that flows through the ports of a distributed port group or uplink port group.

Procedure

1 Locate a distributed port group or an uplink port group in the vSphere Web Client.

   a Select a distributed switch and click the **Related Objects** tab.

   b Click **Distributed Port Groups** to see the list of distributed port groups, or click **Uplink Port Groups** to see the list of uplink port groups.

2 Right-click the port group and select **Edit settings**.

3 Select **Traffic filtering and marking**.

4 If traffic filtering and marking is disabled, enable it from the **Status** drop-down menu.

5 Click **New** to create a new rule, or select a rule and click **Edit** to edit it.

6 In the network traffic rule dialog box, use the Action options to let traffic pass through the ports of the distributed port group or uplink port group, or to restrict it.

7 Specify the kind of traffic that the rule is applicable to.

   To determine if a data flow is in the scope of a rule for marking or filtering, the vSphere distributed switch examines the direction of the traffic, and properties like source and destination, VLAN, next level protocol, infrastructure traffic type, and so on.

   a From the **Traffic direction** drop-down menu, select whether the traffic must be ingress, egress, or both so that the rule recognizes it as matching.

      The direction also influences how you are going to identify the traffic source and destination.

   b By using qualifiers for system data type, Layer 2 packet attributes, and Layer 3 packet attributes set the properties that packets must have to match the rule.

      A qualifier represents a set of matching criteria related to a networking layer. You can match traffic to system data type, Layer 2 traffic properties, and Layer 3 traffic properties. You can use the qualifier for a specific networking layer or can combine qualifiers to match packets more precisely.

      - Use the system traffic qualifier to match packets to the type of virtual infrastructure data that is flowing through the ports of the group . For example, you can select NFS for data transfers to network storage.
      - Use the MAC traffic qualifier to match packets by MAC address, VLAN ID, and next level protocol.

        Locating traffic with a VLAN ID on a distributed port group works with Virtual Guest Tagging (VGT). To match traffic to VLAN ID if Virtual Switch Tagging (VST) is active, use a rule on an uplink port group or uplink port.
      - Use the IP traffic qualifier to match packets by IP version, IP address, and next level protocol and port.

8 In the rule dialog box, click **OK** to save the rule.

Working with Network Traffic Rules on a Distributed Port Group or Uplink Port Group in the vSphere Web Client

Define traffic rules in a distributed port group or uplink port group to introduce a policy for processing traffic related to virtual machines or to physical adapters. You can filter specific traffic or describe its QoS demands.

NOTE You can override the rules of the policy for traffic filtering and marking at port level. See "Working with Network Traffic Rules on a Distributed Port or Uplink Port in the vSphere Web Client," on page 109.

View Traffic Rules on a Distributed Port Group or Uplink Group in the vSphere Web Client

View the traffic rules that form the traffic filtering and marking policy of a distributed port group or uplink port group.

Procedure

1 Locate a distributed port group or an uplink port group in the vSphere Web Client.

   a Select a distributed switch and click the **Related Objects** tab.

   b Click **Distributed Port Groups** to see the list of distributed port groups, or click **Uplink Port Groups** to see the list of uplink port groups.

2 Right-click the port group and select **Edit settings**.

3 Select **Traffic filtering and marking**.

4 If traffic filtering and marking is disabled, enable it from the **Status** drop-down menu.

5 Examine **Action** to see if the rule filters traffic (Allow or Deny) or marks traffic (Tag) with special QoS demands.

6 From the upper list, select the rule for which you want to view the criteria for locating traffic.

   The traffic qualifying parameters of the rule appear in the Traffic Qualifiers list.

Edit a Traffic Rule on a Distributed Port Group or Uplink Port Group in the vSphere Web Client

Create or edit traffic rules, and use their parameters to configure a policy for filtering or marking the traffic on a distributed port group or uplink port group.

Procedure

1 Locate a distributed port group or an uplink port group in the vSphere Web Client.

   a Select a distributed switch and click the **Related Objects** tab.

   b Click **Distributed Port Groups** to see the list of distributed port groups, or click **Uplink Port Groups** to see the list of uplink port groups.

2 Right-click the port group and select **Edit settings**.

3 Select **Traffic filtering and marking**.

4 If traffic filtering and marking is disabled, enable it from the **Status** drop-down menu.

5 Click **New** to create a new rule, or select a rule and click **Edit** to edit it.

What to do next

Name the network traffic rule, and deny, allow, or tag the target traffic.

Change Rule Priorities on a Distributed Port Group or Uplink Port Group in the vSphere Web Client

Reorder the rules that form the traffic filtering and marking policy of a distributed port group or uplink port group to change the sequence of actions for processing traffic.

The vSphere distributed switch applies network traffic rules in a strict order. If a packet already satisfies a rule, the packet might not be passed to the next rule in the policy.

Procedure

1 Locate a distributed port group or an uplink port group in the vSphere Web Client.

   a Select a distributed switch and click the **Related Objects** tab.

   b Click **Distributed Port Groups** to see the list of distributed port groups, or click **Uplink Port Groups** to see the list of uplink port groups.

2 Right-click the port group and select **Edit settings**.

3 Select **Traffic filtering and marking**.

4 If traffic filtering and marking is disabled, enable it from the **Status** drop-down menu.

5 Select a rule and use the arrow buttons to change its priority.

6 Click **OK** to apply the changes.

Delete a Traffic Rule on a Distributed Port Group or Uplink Port Group in the vSphere Web Client

Delete a traffic rule on a distributed port group or uplink port group to stop processing packets flowing to virtual machines or physical adapters in a specific way.

Procedure

1 Locate a distributed port group or an uplink port group in the vSphere Web Client.

   a Select a distributed switch and click the **Related Objects** tab.

   b Click **Distributed Port Groups** to see the list of distributed port groups, or click **Uplink Port Groups** to see the list of uplink port groups.

2 Right-click the port group and select **Edit settings**.

3 Select **Traffic filtering and marking**.

4 If traffic filtering and marking is disabled, enable it from the **Status** drop-down menu.

5 Select the rule and click **Delete**.

6 Click **OK**.

Disable Traffic Filtering and Marking on a Distributed Port Group or Uplink Port Group in the vSphere Web Client

Let traffic flow to virtual machines or physical adapters without additional control related to security or QoS by disabling the traffic filtering and marking policy.

NOTE You can enable and set up the traffic filtering and marking policy on a particular port. See "Enable Traffic Filtering and Marking on a Distributed Port or Uplink Port in the vSphere Web Client," on page 106.

Procedure

1 Locate a distributed port group or an uplink port group in the vSphere Web Client.

   a Select a distributed switch and click the **Related Objects** tab.

   b Click **Distributed Port Groups** to see the list of distributed port groups, or click **Uplink Port Groups** to see the list of uplink port groups.

2 Right-click the port group and select **Edit settings**.

3 Select **Traffic filtering and marking**.

4 From the **Status** drop-down menu, select **Disabled**.

5 Click **OK**.

Traffic Filtering and Marking on a Distributed Port or Uplink Port in the vSphere Web Client

Filter traffic or describe its QoS demands for an individual virtual machine, VMkernel adapter, or physical adapter by configuring the traffic filtering and marking policy on a distributed port or uplink port.

- Enable Traffic Filtering and Marking on a Distributed Port or Uplink Port in the vSphere Web Client on page 106

  Enable the traffic filtering and marking policy on a port to configure traffic security and marking on a virtual machine network adapter, VMkernel adapter, or uplink adapter.

- Mark Traffic on a Distributed Port or Uplink Port in the vSphere Web Client on page 106

  Assign priority tags in a rule for traffic that needs special treatment such as VoIP and streaming video. You can mark the traffic for a virtual machine, VMkernel adapter, or physical adapter with a CoS tag in Layer 2 of the network protocol stack or with a DSCP tag in Layer 3.

- Filter Traffic on a Distributed Port or Uplink Port in the vSphere Web Client on page 108

  By using a rule, permit or stop traffic for securing data flows through a virtual machine, a VMkernel adapter, or a physical adapter.

- Working with Network Traffic Rules on a Distributed Port or Uplink Port in the vSphere Web Client on page 109

  Define traffic rules in a distributed port or uplink port group to introduce a policy for processing traffic related to a virtual machine or to a physical adapter. You can filter specific traffic or describe its QoS demands.

- Disable Traffic Filtering and Marking on a Distributed Port or Uplink Port in the vSphere Web Client on page 111

  Disable the traffic filtering and marking policy on a port to let traffic flow to a virtual machine or a physical adapter without filtering for security or marking for QoS.

Enable Traffic Filtering and Marking on a Distributed Port or Uplink Port in the vSphere Web Client

Enable the traffic filtering and marking policy on a port to configure traffic security and marking on a virtual machine network adapter, VMkernel adapter, or uplink adapter.

Prerequisites

Enable the port-level override option for this policy. See "Edit Advanced Distributed Port Group Settings with the vSphere Web Client," on page 44.

Procedure

1 Browse to a distributed switch in the vSphere Web Client.

2 Navigate to a distributed port or an uplink port.

- To navigate to the distributed ports of the switch, click **Manage > Ports**.
- To navigate to the uplink ports of an uplink port group, click **Related Objects > Uplink Port Groups**, double-click an uplink port group from the list, and select **Ports** on the **Manage** tab.

3 Select a port from the list.

4 Click **Edit distributed port settings**.

5 Select **Traffic filtering and marking**.

6 Select the **Override** check box, and from the **Status** drop-down menu, select **Enabled**.

7 Click **OK**.

What to do next

Set up traffic filtering or marking for the data flowing through the distributed port or through the uplink port. See "Mark Traffic on a Distributed Port or Uplink Port in the vSphere Web Client," on page 106 and "Filter Traffic on a Distributed Port or Uplink Port in the vSphere Web Client," on page 108.

Mark Traffic on a Distributed Port or Uplink Port in the vSphere Web Client

Assign priority tags in a rule for traffic that needs special treatment such as VoIP and streaming video. You can mark the traffic for a virtual machine, VMkernel adapter, or physical adapter with a CoS tag in Layer 2 of the network protocol stack or with a DSCP tag in Layer 3.

Priority tagging is a mechanism to mark traffic that has higher QoS demands. In this way, the network can recognize different classes of traffic. The network devices can handle the traffic from each class according to its priority and requirements.

You can also re-tag traffic to either raise or lower the importance of the flow. By using a low QoS tag, you can restrict data tagged in a guest operating system.

Prerequisites

Enable the port-level override option for this policy. See "Edit Advanced Distributed Port Group Settings with the vSphere Web Client," on page 44.

Procedure

1 Navigate to a distributed port or an uplink port.

- To navigate to the distributed ports of the switch, click **Manage > Ports**.
- To navigate to the uplink ports of an uplink port group, click **Related Objects > Uplink Port Groups**, double-click an uplink port group from the list, and select **Ports** on the **Manage** tab.

2 Select a port from the list.

3 Click **Edit distributed port settings**.

4 If traffic filtering and marking is not enabled at the port level, click **Override**, and from the **Status** drop-down menu, select **Enabled**.

5 Click **New** to create a new rule, or select a rule and click **Edit** to edit it.

You can change a rule inherited from the distributed port group or uplink port group. In this way, the rule becomes unique within the scope of the port.

6 In the network traffic rule dialog box, select the **Tag** option from the **Action** drop-down menu.

7 Set the priority tag for the traffic within the scope of the rule.

Option	Description
CoS value	Mark the traffic matching the rule with a CoS priority tag in network Layer 2. Select **Update CoS tag** and type a value from 0 to 7.
DSCP value	Mark the traffic associated with the rule with a DSCP tag in network Layer 3. Select **Update DSCP value** and type a value from 0 to 63.

8 Specify the kind of traffic that the rule is applicable to.

To determine if a data flow is in the scope of a rule for marking or filtering, the vSphere distributed switch examines the direction of the traffic, and properties like source and destination, VLAN, next level protocol, infrastructure traffic type, and so on.

a From the **Traffic direction** drop-down menu, select whether the traffic must be ingress, egress, or both so that the rule recognizes it as matching.

The direction also influences how you are going to identify the traffic source and destination.

b By using qualifiers for system data type, Layer 2 packet attributes, and Layer 3 packet attributes set the properties that packets must have to match the rule.

A qualifier represents a set of matching criteria related to a networking layer. You can match traffic to system data type, Layer 2 traffic properties, and Layer 3 traffic properties. You can use the qualifier for a specific networking layer or can combine qualifiers to match packets more precisely.

- Use the system traffic qualifier to match packets to the type of virtual infrastructure data that is flowing through the ports of the group . For example, you can select NFS for data transfers to network storage.
- Use the MAC traffic qualifier to match packets by MAC address, VLAN ID, and next level protocol.

  Locating traffic with a VLAN ID on a distributed port group works with Virtual Guest Tagging (VGT). To match traffic to VLAN ID if Virtual Switch Tagging (VST) is active, use a rule on an uplink port group or uplink port.
- Use the IP traffic qualifier to match packets by IP version, IP address, and next level protocol and port.

9 In the rule dialog box, click **OK** to save the rule.

Filter Traffic on a Distributed Port or Uplink Port in the vSphere Web Client

By using a rule, permit or stop traffic for securing data flows through a virtual machine, a VMkernel adapter, or a physical adapter.

Prerequisites

Enable the port-level override option for this policy. See "Edit Advanced Distributed Port Group Settings with the vSphere Web Client," on page 44.

Procedure

1 Navigate to a distributed port or an uplink port.

- To navigate to the distributed ports of the switch, click **Manage > Ports**.
- To navigate to the uplink ports of an uplink port group, click **Related Objects > Uplink Port Groups**, double-click an uplink port group from the list, and select **Ports** on the **Manage** tab.

2 Select a port from the list.

3 Click **Edit distributed port settings**.

4 If traffic filtering and marking is not enabled at the port level, click **Override**, and from the **Status** drop-down menu, select **Enabled**.

5 Click **New** to create a new rule, or select a rule and click **Edit** to edit it.

You can change a rule inherited from the distributed port group or uplink port group. In this way, the rule becomes unique within the scope of the port.

6 In the network traffic rule dialog box, select the **Allow** action to let traffic pass through the distributed port or uplink port, or the **Drop** action to restrict it.

7 Specify the kind of traffic that the rule is applicable to.

To determine if a data flow is in the scope of a rule for marking or filtering, the vSphere distributed switch examines the direction of the traffic, and properties like source and destination, VLAN, next level protocol, infrastructure traffic type, and so on.

a From the **Traffic direction** drop-down menu, select whether the traffic must be ingress, egress, or both so that the rule recognizes it as matching.

The direction also influences how you are going to identify the traffic source and destination.

b By using qualifiers for system data type, Layer 2 packet attributes, and Layer 3 packet attributes set the properties that packets must have to match the rule.

A qualifier represents a set of matching criteria related to a networking layer. You can match traffic to system data type, Layer 2 traffic properties, and Layer 3 traffic properties. You can use the qualifier for a specific networking layer or can combine qualifiers to match packets more precisely.

- Use the system traffic qualifier to match packets to the type of virtual infrastructure data that is flowing through the ports of the group . For example, you can select NFS for data transfers to network storage.
- Use the MAC traffic qualifier to match packets by MAC address, VLAN ID, and next level protocol.

  Locating traffic with a VLAN ID on a distributed port group works with Virtual Guest Tagging (VGT). To match traffic to VLAN ID if Virtual Switch Tagging (VST) is active, use a rule on an uplink port group or uplink port.
- Use the IP traffic qualifier to match packets by IP version, IP address, and next level protocol and port.

8 In the rule dialog box, click **OK** to save the rule.

Working with Network Traffic Rules on a Distributed Port or Uplink Port in the vSphere Web Client

Define traffic rules in a distributed port or uplink port group to introduce a policy for processing traffic related to a virtual machine or to a physical adapter. You can filter specific traffic or describe its QoS demands.

- View Traffic Rules on a Distributed Port or Uplink Port in the vSphere Web Client on page 109

  Review the traffic rules that form the traffic filtering and marking policy of a distributed port or uplink port.

- Edit a Traffic Rule on a Distributed Port or Uplink Port in the vSphere Web Client on page 110

  Create or edit traffic rules, and use their parameters to configure a policy for filtering or marking the traffic on a distributed port or uplink port.

- Change Rule Priorities on a Distributed Port or Uplink Port in the vSphere Web Client on page 110

  Reorder the rules that form the traffic filtering and marking policy of a distributed port or uplink port to change the sequence of actions for analyzing traffic for security and QoS.

- Delete a Traffic Rule on a Distributed Port or Uplink Port in the vSphere Web Client on page 111

  Delete a traffic rule on a distributed port or uplink port to stop filtering or marking certain type of packets that are flowing to a virtual machine or a physical adapter.

View Traffic Rules on a Distributed Port or Uplink Port in the vSphere Web Client

Review the traffic rules that form the traffic filtering and marking policy of a distributed port or uplink port.

Prerequisites

Enable the port-level override option for this policy. See "Edit Advanced Distributed Port Group Settings with the vSphere Web Client," on page 44.

Procedure

1 Navigate to a distributed port or an uplink port.

   - To navigate to the distributed ports of the switch, click **Manage > Ports**.
   - To navigate to the uplink ports of an uplink port group, click **Related Objects > Uplink Port Groups**, double-click an uplink port group from the list, and select **Ports** on the **Manage** tab.

2 Select a port from the list.

3 Click **Edit distributed port settings**.

4 Select **Traffic filtering and marking**.

5 If traffic filtering and marking is not enabled at the port level, click **Override**, and from the **Status** drop-down menu, select **Enabled**.

6 Examine **Action** to see if the rule filters traffic (Allow or Deny) or marks traffic (Tag) with special QoS demands.

7 From the upper list, select the rule for which you want to view the criteria for locating traffic.

   The traffic qualifying parameters of the rule appear in the Traffic Qualifiers list.

Edit a Traffic Rule on a Distributed Port or Uplink Port in the vSphere Web Client

Create or edit traffic rules, and use their parameters to configure a policy for filtering or marking the traffic on a distributed port or uplink port.

Prerequisites

Enable the port-level override option for this policy. See "Edit Advanced Distributed Port Group Settings with the vSphere Web Client," on page 44.

Procedure

1 Navigate to a distributed port or an uplink port.

- To navigate to the distributed ports of the switch, click **Manage > Ports**.
- To navigate to the uplink ports of an uplink port group, click **Related Objects > Uplink Port Groups**, double-click an uplink port group from the list, and select **Ports** on the **Manage** tab.

2 Select a port from the list.

3 Click **Edit distributed port settings**.

4 Select **Traffic filtering and marking**.

5 If traffic filtering and marking is not enabled at the port level, click **Override**, and from the **Status** drop-down menu, select **Enabled**.

6 Click **New** to create a new rule, or select a rule and click **Edit** to edit it.

You can change a rule inherited from the distributed port group or uplink port group. In this way, the rule becomes unique within the scope of the port.

What to do next

Name the network traffic rule, and deny, allow, or tag the target traffic.

Change Rule Priorities on a Distributed Port or Uplink Port in the vSphere Web Client

Reorder the rules that form the traffic filtering and marking policy of a distributed port or uplink port to change the sequence of actions for analyzing traffic for security and QoS.

The vSphere distributed switch applies network traffic rules in a strict order. If a packet already satisfies a rule, the packet might not be passed to the next rule in the policy.

Prerequisites

Enable the port-level override option for this policy. See "Edit Advanced Distributed Port Group Settings with the vSphere Web Client," on page 44.

Procedure

1 Navigate to a distributed port or an uplink port.

- To navigate to the distributed ports of the switch, click **Manage > Ports**.
- To navigate to the uplink ports of an uplink port group, click **Related Objects > Uplink Port Groups**, double-click an uplink port group from the list, and select **Ports** on the **Manage** tab.

2 Select a port from the list.

3 Click **Edit distributed port settings**.

4 Select **Traffic filtering and marking**.

5 If traffic filtering and marking is not enabled at the port level, click **Override**, and from the **Status** drop-down menu, select **Enabled**.

6 Select a rule and use the arrow buttons to change its priority.

7 Click **OK** to apply the changes.

Delete a Traffic Rule on a Distributed Port or Uplink Port in the vSphere Web Client

Delete a traffic rule on a distributed port or uplink port to stop filtering or marking certain type of packets that are flowing to a virtual machine or a physical adapter.

Prerequisites

Enable the port-level override option for this policy. See "Edit Advanced Distributed Port Group Settings with the vSphere Web Client," on page 44.

Procedure

1 Navigate to a distributed port or an uplink port.

- To navigate to the distributed ports of the switch, click **Manage > Ports**.
- To navigate to the uplink ports of an uplink port group, click **Related Objects > Uplink Port Groups**, double-click an uplink port group from the list, and select **Ports** on the **Manage** tab.

2 Select a port from the list.

3 Click **Edit distributed port settings**.

4 Select **Traffic filtering and marking**.

5 If traffic filtering and marking is not enabled at the port level, click **Override**, and from the **Status** drop-down menu, select **Enabled**.

6 Select the rule and click **Delete**.

7 Click **OK**.

Disable Traffic Filtering and Marking on a Distributed Port or Uplink Port in the vSphere Web Client

Disable the traffic filtering and marking policy on a port to let traffic flow to a virtual machine or a physical adapter without filtering for security or marking for QoS.

Prerequisites

Enable the port-level override option for this policy. See "Edit Advanced Distributed Port Group Settings with the vSphere Web Client," on page 44.

Procedure

1 Navigate to a distributed port or an uplink port.

- To navigate to the distributed ports of the switch, click **Manage > Ports**.
- To navigate to the uplink ports of an uplink port group, click **Related Objects > Uplink Port Groups**, double-click an uplink port group from the list, and select **Ports** on the **Manage** tab.

2 Select a port from the list.

3 Click **Edit distributed port settings**.

4 Select **Traffic filtering and marking**.

5 Click **Override**, and from the **Status** drop-down menu, select **Disabled**.

6 Click **OK**.

Qualifying Traffic for Filtering and Marking

The traffic that you want to filter or want to mark with QoS tags can be matched to the type of carried infrastructure data, such as data for storage, vCenter Server management, and so on, and to Layer 2 and Layer 3 properties .

To match the traffic in the scope of the rule more precisely, you can combine criteria for system data type, Layer 2 header, and Layer 3 header.

System Traffic Qualifier

By using the system traffic qualifier in a rule for port group or port, you can determine whether certain system data traffic must be marked with a QoS tag, allowed, or dropped.

System Traffic Type

You can select the type of traffic through the ports of the group that carries system data, that is, traffic for management from vCenter Server , storage, VMware vSphere® vMotion®, and vSphere Fault Tolerance. You can mark or filter only a specific traffic type, or for all system data traffic except for a infrastructure feature. For example, you can mark with a QoS value or filter the traffic for management from vCenter Server , storage and vMotion, but not the traffic carrying the Fault Tolerance data.

MAC Traffic Qualifier

By using the MAC traffic qualifier in a rule, you can define matching criteria for the Layer 2 (Data Link Layer) properties of packets such as MAC address, VLAN ID, and next level protocol that consumes the frame payload.

Protocol Type

The **Protocol type** attribute of the MAC traffic qualifier corresponds to the EtherType field in Ethernet frames. EtherType represents the type of next level protocol that is going to consume the payload of the frame.

You can select a protocol from the drop-down menu or type its hexadecimal number. For example, to capture traffic for the Link Layer Discovery Protocol (LLDP) protocol, type **88CC**.

VLAN ID

You can use the VLAN ID attribute of the MAC traffic qualifier to mark or filter traffic in a particular VLAN.

NOTE The VLAN ID qualifier on a distributed port group works with Virtual Guest Tagging (VGT).

If a flow is tagged with a VLAN ID through Virtual Switch Tagging (VST), it cannot be located by using this ID in a rule on a distributed port group or distributed port. The reason is that the distributed switch checks the rule conditions, including the VLAN ID, after the switch has already untagged the traffic. In this case, to match traffic by VLAN ID successfully, you must use a rule on an uplink port group or uplink port.

Source Address

By using the Source Address group of attributes, you can match packets by the source MAC address or network.

You can use a comparison operator to mark or filter packets that have or do not have the specified source address or network.

You can match the traffic source in several ways.

Table 6-1. Patterns for Filtering or Marking Traffic by MAC Source Address

Parameters to Match Traffic Source Address	Comparison Operator	Networking Argument Format
MAC address	**is** or **is not**	Type the MAC address for matching. Use colons to separate the octets in it.
MAC network	**matches** or **does not match**	Type the lowest address in the network and a wildcard mask. Set zeroes at the positions of the network bits, and ones for the host part.

For example, for a MAC network with prefix 05:50:56 that is 23 bits long, set the address as **`00:50:56:00:00:00`** and mask as **`00:00:01:ff:ff:ff`**.

Destination Address

By using the Destination Address group of attributes, you can match packets to their destination address. The MAC destination address options have the same format as those for the source address.

Comparison Operators

To match traffic in a MAC qualifier more closely to your needs, you can use affirmative comparison or negation. You can use operators such that all packets except the ones with certain attributes fall in the scope of a rule.

IP Traffic Qualifier

By using the IP traffic qualifier in a rule, you can define criteria for matching traffic to the Layer 3 (Network Layer) properties such as IP version, IP address, next level protocol, and port.

Protocol

The **Protocol** attribute of the IP traffic qualifier represents the next level protocol consuming the payload of the packet. You can select a protocol from the drop-down menu or type its decimal number according to RFC 1700.

For the TCP and UDP protocols, you can also match traffic by source and destination ports.

Source Port

By using the Source port attribute, you can match TCP or UDP packets by the source port. Consider the traffic direction when matching traffic to a source port.

Destination Port

By using the Destination port attribute, you can match TCP or UDP packets by the destination port. Consider the traffic direction when matching traffic to a destination port.

Source Address

By using the Source Address attribute, you can match packets by source address or subnet. Consider the traffic direction when matching traffic to a source address or network.

You can match traffic source in several ways.

Table 6-2. Patterns for Filtering or Marking Traffic by IP Source Address

Parameters to Match Traffic Source Address	Comparison Operator	Networking Argument Format
IP version	**any**	Select the IP version from the drop-down menu.
IP address	**is** or **is not**	Type the IP address that you want to match.
IP subnet	**matches** or **does not match**	Type the lowest address in the subnet and the bit length of the subnet prefix.

Destination Address

Use the Destination Address to match packets by IP address, subnet, or IP version. The destination address has the same format as the one for the source.

Comparison Operators

To match traffic in an IP qualifier more closely to your needs, you can use affirmative comparison or negation. You can define that all packets fall in the scope of a rule except packets with certain attributes.

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 in the vSphere Web Client

You can configure various distributed port group policies.

Procedure

1 Browse to a distributed switch in the vSphere Web Client navigator.

2 Right-click the distributed switch in the object navigator and select **Manage Distributed Port Groups**.

3 Select the **Miscellaneous** check box and click **Next**.

4 Select a distributed port group to edit and click **Next**.

5 Use the **Block all ports** drop-down menu to select **Yes** or **No** and click **Next**.

Selecting Yes shuts down all ports in the port group. This might disrupt the normal network operations of the hosts or virtual machines using the ports.

6 Review your settings and click **Finish**.

Use the **Back** button to change any settings.

Edit Distributed Port or Uplink Port Blocking Policies with the vSphere Web Client

You can configure distributed port or uplink port blocking policies.

Prerequisites

To override the traffic shaping policy at the port level, enable the port-level overrides. See "Edit Advanced Distributed Port Group Settings with the vSphere Web Client," on page 44

Procedure

1 Browse to a distributed switch in the vSphere Web Client navigator.

2 Click the **Manage** tab, and select **Ports**.

3 Select a port from the list.

4 Click **Edit distributed port settings**.

5 In the **Miscellaneous** section, select the **Block Port Override** check box and choose **Yes** or **No** from the drop-down menu.

Yes shuts down all ports in the port group. This might disrupt the normal network operations of the hosts or virtual machines using the ports.

6 Click **OK**.

Manage Policies for Multiple Port Groups on a vSphere Distributed Switch in the vSphere Web Client

You can modify networking policies for multiple port groups on a vSphere Distributed Switch.

Prerequisites

Create a vSphere Distributed Switch with one or more port groups.

Procedure

1 Browse to a distributed switch in the vSphere Web Client.

2 Right-click the distributed switch, and select **Manage Distributed Port Groups**.

3 On the Select port group policies page, select the check box next to the policy categories to modify and click **Next**.

Option	Description
Security	Set MAC address changes, forged transmits, and promiscuous mode for the selected port groups.
Traffic shaping	Set the average bandwidth, peak bandwidth, and burst size for inbound and outbound traffic on the selected port groups.
VLAN	Configure how the selected port groups connect to physical VLANs.
Teaming and failover	Set load balancing, failover detection, switch notification, and failover order for the selected port groups.
Resource allocation	Set network resource pool association for the selected port groups. This option is available for vSphere Distributed Switch version 5.0. and later.
Monitoring	Enable or disable NetFlow on the selected port groups. This option is available for vSphere Distributed Switch version 5.0.0 and later.
Traffic filtering and marking	Configure policy for filtering (allow or drop) and for marking certain types of traffic through the ports of selected port groups. This option is available for vSphere Distributed Switch version 5.5 and later .
Miscellaneous	Enable or disable port blocking on the selected port groups.

4 On the Select port groups page, select the distributed port group(s) to edit and click **Next**.

5 (Optional) On the Security page, use the drop-down menus to edit the security exceptions and click **Next**.

Option	Description
Promiscuous Mode	■ **Reject**. Placing a guest adapter in promiscuous mode has no effect on which frames are received by the adapter. ■ **Accept**. Placing a guest adapter in promiscuous mode causes it to detect all frames passed on the vSphere Distributed Switch that are allowed under the VLAN policy for the port group that the adapter is connected to.
MAC Address Changes	■ **Reject**. If set to **Reject** and the guest operating system changes the MAC address of the adapter to anything other than what is in the `.vmx` configuration file, all inbound frames are dropped. If the Guest OS changes the MAC address back to match the MAC address in the `.vmx` configuration file, inbound frames are passed again. ■ **Accept**. Changing the MAC address from the Guest OS has the intended effect. Frames to the new MAC address are received.
Forged Transmits	■ **Reject**. Any outbound frame with a source MAC address that is different from the one currently set on the adapter are dropped. ■ **Accept**. No filtering is performed and all outbound frames are passed.

6 (Optional) On the Traffic shaping page, use the drop-down menus to enable or disable Ingress or Egress traffic shaping and click **Next**.

Option	Description
Status	If you enable either **Ingress Traffic Shaping** or **Egress Traffic Shaping**, you are setting limits on the amount of networking bandwidth allocated for each VMkernel adapter or virtual network adapter associated with this port group. If you disable the policy, services have a free, clear connection to the physical network by default.
Average Bandwidth	Establishes the number of bits per second to allow across a port, averaged over time, that is, the allowed average load.
Peak Bandwidth	The maximum number of bits per second to allow across a port when it is sending or receiving a burst of traffic. This maximum number tops the bandwidth used by a port whenever it is using its burst bonus.
Burst Size	The maximum number of bytes to allow in a burst. If this parameter is set, a port might gain a burst bonus when it does not use all its allocated bandwidth. Whenever the port needs more bandwidth than specified by **Average Bandwidth**, it might be allowed to transmit data at a higher speed if a burst bonus is available. This parameter tops the number of bytes that can be accumulated in the burst bonus and transferred at a higher speed.

7 (Optional) On the VLAN page, use the drop-down menus to edit the VLAN policy and click **Next**.

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

8 (Optional) On the Teaming and failover page, use the drop-down menus to edit the settings and click **Next**.

Option	Description
Load Balancing	IP-based teaming requires that the physical switch be configured with ether channel. For all other options, ether channel should be disabled. Select how to choose an uplink. ■ **Route based on the originating virtual port**. Choose an uplink based on the virtual port where the traffic entered the distributed switch. ■ **Route based on IP hash**. Choose an uplink based on a hash of the source and destination IP addresses of each packet. For non-IP packets, whatever is at those offsets is used to compute the hash. ■ **Route based on source MAC hash**. Choose an uplink based on a hash of the source Ethernet. ■ **Route based on physical NIC load**. Choose an uplink based on the current loads of physical NICs. ■ **Use explicit failover order**. Always use the highest order uplink, from the list of Active adapters, which passes failover detection criteria.
Network Failover Detection	Select the method to use for failover detection. ■ **Link Status only**. Relies solely on the link status that the network adapter provides. This option detects failures, such as cable pulls and physical switch power failures, but not configuration errors, such as a physical switch port being blocked by spanning tree or that is misconfigured to the wrong VLAN or cable pulls on the other side of a physical switch. ■ **Beacon Probing**. Sends out and listens for beacon probes on all NICs in the team and uses this information, in addition to link status, to determine link failure. Do not use beacon probing with IP-hash load balancing.
Notify Switches	Select **Yes** or **No** to notify switches in the case of failover. Do not use this option when the virtual machines using the port group are using Microsoft Network Load Balancing in unicast mode. If you select **Yes**, whenever a virtual NIC is connected to the distributed 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 out over the network to update the lookup tables on physical switches. Use this process for the lowest latency of failover occurrences and migrations with vMotion.
Failback	Select **Yes** or **No** to disable or enable failback. This option determines how a physical adapter is returned to active duty after recovering from a failure. ■ **Yes** (default). The adapter is returned to active duty immediately upon recovery, displacing the standby adapter that took over its slot, if any. ■ **No**. A failed adapter is left inactive even after recovery until another currently active adapter fails, requiring its replacement.
Failover Order	Select how to distribute the work load for uplinks. To use some uplinks but reserve others in case the uplinks in use fail, set this condition by moving them into different groups. ■ **Active Uplinks**. Continue to use the uplink when the network adapter connectivity is up and active. ■ **Standby Uplinks** . Use this uplink if one of the active adapter's connectivity is down. When using IP-hash load balancing, do not configure standby uplinks. ■ **Unused Uplinks** . Do not use this uplink.

9 (Optional) On the Resource allocation page, use the network resource pool drop-down menu to add or remove resource allocations and click **Next**.

10 (Optional) On the Monitoring page, use the drop-menu to enable or disable NetFlow and click **Next**.

Option	Description
Disabled	NetFlow is disabled on the distributed port group.
Enabled	NetFlow is enabled on the distributed port group. You can configure NetFlow settings at the vSphere Distributed Switch level.

11 (Optional) On the Traffic filtering and marking page, enable or disable traffic filtering and marking from **Status** drop-down menu, configure traffic rules for filtering or marking specific data flows, and click **Next**.

You can set the following attributes of a rule determining the target traffic and the action on it:

Option	Description
Name	Name of the rule
Action	■ **Allow**. Grant access to traffic of a certain type. ■ **Drop**. Deny access to traffic of a certain type. ■ **Tag.** Classify traffic in terms of QoS by inserting or retagging traffic with a CoS and DSCP tag.
Traffic direction	Set whether the rule is for incoming, outgoing or incoming and outgoing traffic. The direction also influences how you are going to identify the traffic source and destination.
System traffic qualifier	Indicate that the rule scopes over system traffic and set the type of infrastructure protocol to apply the rule on. For example, mark with a priority tag the traffic for management from vCenter Server.

Option	Description
MAC qualifier	Qualify the traffic for the rule by Layer 2 header. ■ **Protocol type**. Set the next level protocol (IPv4, IPv6, etc.) consuming the payload. This attribute corresponds to the EtherType field in Ethernet frames. You can select a protocol from the drop-down menu or type its hexadecimal number For example, to locate traffic for the Link Layer Discovery Protocol (LLDP) protocol, type **88CC**. ■ **VLAN ID**. Locate traffic by VLAN. The VLAN ID qualifier on a distributed port group works with Virtual Guest Tagging (VGT). If you have a flow tagged with a VLAN ID through Virtual Switch Tagging (VST), you cannot locate the flow by this ID in a distributed port group rule. The reason is that the distributed switch checks the rule conditions, including the VLAN ID, after the switch has already untagged the traffic. To match successfully traffic to a VLAN ID, use a rule for an uplink port group or uplink port. ■ **Source Filter**. Set a single MAC address or a MAC network to match packets by source address. For a MAC network you enter the lowest address in the network and a wildcard mask. The mask contains zeroes at the positions of the network bits, and ones for the host part. For example, for a MAC network with prefix 05:50:56 that is 23 bits long, set the address as **00:50:56:00:00:00** and the mask as **00:00:01:ff:ff:ff**. ■ **Destination Filter**. Set a single MAC address or a MAC network to match packets by destination address. The MAC destination address supports the same format as the source address.
IP qualifier	Qualify the traffic for the rule by Layer 3 header. ■ **Protocol**. Set the next level protocol (TCP, UDP, etc.) consuming the payload. You can select a protocol from the drop-down menu or type its decimal number according to *RFC 1700, Assigned Numbers*. For TCP and UDP protocol, you can also set source and destination port. ■ **Source port**. Match TCP or UDP packets to a source port. Consider the direction of the traffic that is within the scope of the rule when determining the source port to match packets to. ■ **Destination port**. Match TCP or UDP packets by the source port. Consider the direction of the traffic that is within the scope of the rule when determining the destination port to match packets to. ■ **Source Filter**. Set the IP version, a single IP address or a subnet to match packets by source address. For a subnet you enter the lowest address and the bit length of the prefix. ■ **Destination Filter**. Set the IP version, a single IP address or a subnet to match packets by source address. The IP destination address supports the same format as the source address.

12 (Optional) On the Miscellaneous page, select **Yes** or **No** from the drop-down menu and click **Next**.

Select**Yes** to shut down all ports in the port group. This shutdown might disrupt the normal network operations of the hosts or virtual machines using the ports.

13 Review your settings on the Ready to complete page and click **Finish**.

Use the **Back** button to change any settings.

Managing Network Resources 7

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

This chapter includes the following topics:

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 shares 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 with the vSphere Web Client on page 122

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

- Create a Network Resource Pool with the vSphere Web Client on page 122

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

- Add or Remove Distributed Port Groups from a Network Resource Pool with the vSphere Web Client on page 123

  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 with the vSphere Web Client on page 124

  You can change both system and user-defined network resource pool settings to change the priority network traffic for that network resource pool.

- Delete a User-Defined Network Resource Pool with the vSphere Web Client on page 124

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

Enable Network I/O Control on a vSphere Distributed Switch with the vSphere Web Client

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 Browse to a distributed switch in the vSphere Web Client navigator.

2 Click the **Manage** tab, and select **Settings > Properties**.

3 Click **Edit**.

4 Select to **Enable** or **Disable** network I/O control from the **Network I/O Control** drop-down menu.

5 Click **OK**.

Create a Network Resource Pool with the vSphere Web Client

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 Browse to a distributed switch in the vSphere Web Client navigator.

2 Click the **Manage** tab, and click **Resource Allocation**.

3 Click **New**.

4 Type a **Name** for the network resource pool or accept the generated name.

5 (Optional) Type a **Description** for the network resource pool.

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

7 Select the **Physical adapter shares** for the network resource pool from the drop-down menu.

Option	Description
Low	Sets the shares for this resource pool to 25.
Normal	Sets the shares for this resource pool to 50.
High	Sets the shares for this resource pool to 100.
Custom	Aspecific number of shares, from 1 to 100, for this network resource pool.

8 (Optional) Select the **QoS tag** for the network resource pool.

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

9 Click **OK**.

The new resource pool appears in the **User-defined network resource pools** section.

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 with the vSphere Web Client

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 user-defined network resource pools on the vSphere distributed switch.

Procedure

1 Browse to a distributed switch in the vSphere Web Client.

2 Right-click the distributed switchin the navitgator and select **Manage Distributed Port Groups**.

3 On the **Select port group policies** page, select the **Resource allocation** check box and click **Next**.

4 On the **Select port groups** page, select port groups to edit and click **Next**.

5 On the **Configure policies - Resource allocation** page, add or remove the distributed switch from the network resource pool and click **Next**.

- To **add** the distributed port group to a resource pool, select a user-defined resource pool from the **Network resource pool** drop-down menu.
- To **remove** the distributed port group from a resource pool, select **default** from the **Network resource pool** drop-down menu.

NOTE If there are no user-defined network resource pools on the distributed switch, you will only see **default** in the drop-down menu.

6 Review your settings on the **Ready to Complete** page and click **Finish**.

Use the **Back** button to change your selections.

Edit Network Resource Pool Settings with the vSphere Web Client

You can change both system and user-defined network resource pool settings to change the priority network traffic for that network resource pool.

Procedure

1 Browse to a distributed switch in the vSphere Web Client navigator.

2 Click the **Manage** tab, and click **Resource Allocation**.

3 Select a network resource pool from the list and click **Edit**.

4 Edit the **Host limit** for the network resource pool in megabits per second or select **Unlimited**.

5 Select an option for the network resource pool from the **Physical adapter shares** drop-down menu.

Option	Description
Low	Sets the shares for this resource pool to 25.
Normal	Sets the shares for this resource pool to 50.
High	Sets the shares for this resource pool to 100.
Custom	A specific number of shares, from 1 to 100, for this network resource pool.

6 (Optional) Select a **QoS tag** for the network resource pool.

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 User-Defined Network Resource Pool with the vSphere Web Client

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

NOTE You cannot delete a system network resource pool.

Prerequisites

Remove all distributed port groups from the network resource pool.

Procedure

1 Browse to a distributed switch in the vSphere Web Client navigator.

2 Click the **Manage** tab, and click **Resource Allocation**.

3 Select a user-defined resource allocation pool and click **Remove**.

4 Click **Yes** to delete the resource pool.

TCP Segmentation Offload and Jumbo Frames

Using TCP Segmentation Offload (TSO) in a VMkernel network adapter and virtual machines, and jumbo frames on a vSphere distributed switch or vSphere standard switch, improves the network performance in virtual machines and infrastructure workloads.

Enabling TSO

Use TCP Segmentation Offload (TSO) in VMkernel network adapters and virtual machines to improve the network performance.

TSO on the transmission path of physical and virtual machine NICs improves the performance of ESX/ESXi hosts by reducing the overhead of the CPU for TCP/IP network operations. When TSO is enabled, the NIC divides larger data chunks into TCP segments instead of the CPU. The host can use more CPU cycles to run applications.

Enable TSO Support for a Virtual Machine by Using the vSphere Web Client

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

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.

Prerequisites

To use TSO, verify that the virtual machine runs one of the following guest operating systems:

- Microsoft Windows Server 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)

Procedure

1. Locate the virtual machine in the vSphere Web Client.
   a. Select a datacenter, folder, cluster, resource pool, or host and click the **Related Objects** tab.
   b. Click **Virtual Machines** and select the virtual machine from the list.
2. On the **Manage** tab of the virtual machine, select **Settings > VM Hardware**.
3. Click **Edit** and select the **Virtual Hardware** tab.
4. Expand the network adapter section and write down the network settings and MAC address for the network adapter.
5. Click **Remove** to remove the network adapter from the virtual machine.
6. From the **New device** drop-down menu, select **Network** and click **Add**.
7. From the **Adapter Type** drop-down menu, select **VMXNET 2 (Enhanced)** or **VMXNET 3**.
8. Set the network setting and MAC address used by the old network adapter.
9. Click **OK**.

What to do next

If the virtual machine is not set to upgrade VMware Tools each time it is powered on, you must upgrade VMware Tools manually.

Enable TSO Support for a VMkernel Network Adapter

When using TSO on a VMkernel network adapter, the physical NIC splits large data chunks into TCP segments instead of the CPU. As a result, the CPU has more cycles for applications. By default, a host is configured to use hardware TSO if its NICs support it.

If TSO becomes disabled for a particular VMkernel adapter, the only way you can re-enable TSO is by deleting the adapter and recreating it with TSO enabled.

Prerequisites

Verify that one of the following requirements are satisfied:

- The VMkernel adapter does not handle traffic related to host networking services, such as iSCSI, vSphere vMotion, and so on.
- The networking services can send traffic by using an alternative VMkernel adapter.

Procedure

1. In the vSphere Web Client, navigate to the host.
2. Under **Manage**, select **Networking** and then select **VMkernel adapters**.
3. Select a VMkernel adapter from the list, and note its settings.

   Later you should set the same settings to the newly added VMkernel adapter.
4. Click **Remove**.
5. In the confirmation dialog, click **Analyze impact**, verify that no networking services are impacted, and click **OK**.
6. In the VMkernel adapters list, click **Add host networking**.
7. On the Add connection type page, select **VMkernel Network Adapter** and click **Next**.
8. On the Select target device page, assign the VMkernel adapter to a standard switch or distributed port group.
9. On the Port properties page, configure the IP version and networking services related to the old adapter, and click **Next**.
10. On the IPv4 Settings and IPv6 Settings pages, configure the IP settings used by the old adapter.
11. Review your settings and click **Finish**.

Enabling Jumbo Frames

Jumbo frames let ESXi hosts send larger frames out onto the physical network. The network must support jumbo frames end-to-end that includes physical network adapters, physical switches, and storage devices.

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

You can enable jumbo frames on a vSphere distributed switch or vSphere standard switch by changing the maximum transmission unit (MTU) to a value greater than 1500 bytes. 9000 bytes is the maximum frame size that you can configure.

Enable Jumbo Frames on a vSphere Distributed Switch with the vSphere Web Client

Enable jumbo frames for the entire traffic that passes through a vSphere distributed switch.

Procedure

1 Browse to a distributed switch in the vSphere Web Client navigator.

2 Click the **Manage** tab, and click **Settings > Properties**.

3 Click **Edit**.

4 Click **Advanced** and set the **MTU** property to a value greater than 1500 bytes.

You cannot set the MTU size to a value greater than 9000 bytes.

5 Click **OK**.

Enable Jumbo Frames on a vSphere Standard Switch with the vSphere Web Client

Enable jumbo frames for all traffic through a vSphere standard switch on a host.

Procedure

1 In the vSphere Web Client, navigate to the host.

2 On the **Manage** tab, click **Networking**, and select **Virtual switches**.

3 Select a standard switch from the virtual switch table and click **Edit settings**.

4 In the **Properties** section, set the **MTU** property to a value greater than 1500 bytes.

You can increase the MTU size up to 9000 bytes.

5 Click **OK**.

Enable Jumbo Frames for a VMkernel Adapter in the vSphere Web Client

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

Procedure

1 In the vSphere Web Client, navigate to the host.

2 Under **Manage**, select **Networking** and then select **VMkernel adapters**.

3 Select a VMkernel adapter from the adapter table.

The properties of the adapter appear.

4 Click the name of the VMkernel adapter.

5 Click **Edit**.

6 Select **NIC settings** and set the **MTU** property to a value greater than 1500.

You can increase the MTU size up to 9000 bytes.

7 Click **OK**.

Enable Jumbo Frame Support on a Virtual Machine with the vSphere Web Client

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

Procedure

1 Locate the virtual machine in the vSphere Web Client.

   a Select a datacenter, folder, cluster, resource pool, or host and click the **Related Objects** tab.

   b Click **Virtual Machines** and select the virtual machine from the list.

2 On the **Manage** tab of the virtual machine, select **Settings > VM Hardware**.

3 Click **Edit** and select the **Virtual Hardware** tab.

4 Click the **Virtual Hardware** section, and expand the network adapter section. 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 From the **New device** drop-down menu, select **Network** and click **Add**.

7 From the **Adapter Type** drop-down menu, select **VMXNET 2 (Enhanced)** or **VMXNET 3**.

8 Set the network settings to the ones recorded for the old network adapter.

9 Set the **MAC Address** to **Manual**, and type the MAC address that the old network adapter was using.

10 Click **OK**.

What to do next

- Check that the enhanced VMXNET adapter is connected to a standard switch or to a distributed switch with jumbo frames enabled.
- Inside the guest operating system, configure the network adapter to allow jumbo frames. See the documentation of your guest operating system.
- 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 following command depending on the host version:

ESX/ESXi Version	Command
ESX/ESXi 4.x	`vicfg-advcfg --set true VMkernel.Boot.netNetQueueEnable`
ESXi 5.x	`esxcli system settings kernel set --setting="netNetqueueEnabled" --value="TRUE"`

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 following command depending on the host version:

ESX/ESXi version	Command
ESX/ESXi 4.x	`vicfg-advcfg --set false VMkernel.Boot.netNetQueueEnable`
ESXi 5.x	`esxcli system settings kernel set --setting="netNetqueueEnabled" --value="FALSE"`

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

Cisco Unified Computing Systems (UCS) through Cisco Virtual Machine Fabric Extender (VM-FEX) distributed switches support the following features for migration and resource management of virtual machines which use DirectPath I/O:

- 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.

- Enable Passthrough for a Network Device on a Host in the vSphere Web Client on page 130

  Passthrough devices provide the means to use resources efficiently and improve performance of your environment. You can enable DirectPath I/O passthrough for a network device on a host.

- Configure a PCI Device on a Virtual Machine with the vSphere Web Client on page 131

  Passthrough devices provide the means to more efficiently use resources and improve performance in your environment. You can configure a passthrough PCI device on a virtual machine in the vSphere Web Client.

- Enable DirectPath I/O with vMotion on a Virtual Machine with the vSphere Web Client on page 131

  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.

Enable Passthrough for a Network Device on a Host in the vSphere Web Client

Passthrough devices provide the means to use resources efficiently and improve performance of your environment. You can enable DirectPath I/O passthrough for a network device on a host.

CAUTION If your ESXi host is configured to boot from a USB device or an SD card attached to a USB channel, make sure that you do not enable DirectPath I/O passthrough for the USB controller. Passing through a USB controller on an ESXi host that boots from a USB device or SD card might put the host in a state where its configuration cannot be persisted.

Procedure

1. Browse to a host in the vSphere Web Client navigator.
2. Click the **Manage** tab, click **Settings**.
3. In the Hardware section, click **PCI Devices**.
4. To enable DirectPath I/O passthrough for a PCI network device on the host, click **Edit**.

   A list of available passthrough devices appears.

Icon	Description
green icon	A device is active and can be enabled.
orange icon	The state of the device has changed, and you must reboot the host before you can use the device.

5. Select the network device to be used for passthrough and click **OK**.

   The selected PCI device appears in the table. Device information is displayed at the bottom of the screen.

6. Reboot the host to make the PCI network device available for use.

Configure a PCI Device on a Virtual Machine with the vSphere Web Client

Passthrough devices provide the means to more efficiently use resources and improve performance in your environment. You can configure a passthrough PCI device on a virtual machine in the vSphere Web Client.

Prerequisites

Verify that a passthrough networking device is configured on the host of the virtual machine. See "Enable Passthrough for a Network Device on a Host in the vSphere Web Client," on page 130.

Procedure

1 Locate the virtual machine in the vSphere Web Client.

   a Select a datacenter, folder, cluster, resource pool, or host and click the **Related Objects** tab.

   b Click **Virtual Machines** and select the virtual machine from the list.

2 Power off the virtual machine.

3 On the **Manage** tab of the virtual machine, select **Settings > VM Hardware**.

4 Click **Edit** and select the **Virtual Hardware** tab.

5 Expand the **Memory** section, and set the **Limit** to **Unlimited**.

6 From the **New device** drop-down menu select **PCI Device** and click **Add**.

7 From the **New PCI device** drop-down menu select the passthrough device to use, and click **OK**.

8 Power on the virtual machine.

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 with the vSphere Web Client

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 documentation at the CIsco Web site.http://www.cisco.com/go/unifiedcomputing/b-series-doc

Procedure

1 Locate the virtual machine in the vSphere Web Client.

   a Select a datacenter, folder, cluster, resource pool, or host and click the **Related Objects** tab.

   b Click **Virtual Machines** and select the virtual machine from the list.

2 Power off the virtual machine.

3 On the **Manage** tab of the virtual machine, select **Settings > VM Hardware**.

4 Click **Edit** and select the **Virtual Hardware** tab.

5 Expand the **Memory** section, and set the **Limit** to **Unlimited**.

6 Expand the **Network adapter** section to configure a passthrough device.

7 Select a port profile with high performance enabled from the network drop-down menu and click **OK**.

8 Power on the virtual machine.

Single Root I/O Virtualization (SR-IOV)

vSphere 5.1 and later releases support Single Root I/O Virtualization (SR-IOV). You can use SR-IOV for networking of virtual machines that are latency sensitive or require more CPU resources.

Overview of SR-IOV

SR-IOV is a specification that allows a single Peripheral Component Interconnect Express (PCIe) physical device under a single root port to appear as multiple separate physical devices to the hypervisor or the guest operating system.

SR-IOV uses physical functions (PFs) and virtual functions (VFs) to manage global functions for the SR-IOV devices. PFs are full PCIe functions that are capable of configuring and managing the SR-IOV functionality. It is possible to configure or control PCIe devices using PFs, and the PF has full ability to move data in and out of the device. VFs are lightweight PCIe functions that support data flowing but have a restricted set of configuration resources.

The number of virtual functions provided to the hypervisor or the guest operating system depends on the device. SR-IOV enabled PCIe devices require appropriate BIOS and hardware support, as well as SR-IOV support in the guest operating system driver or hypervisor instance. See "SR-IOV Support," on page 132.

Using SR-IOV in vSphere

In vSphere, a virtual machine can use an SR-IOV virtual function for networking. The virtual machine and the physical adapter exchange data directly without using the VMkernel as an intermediary. Bypassing the VMkernel for networking reduces latency and improves CPU efficiency.

In vSphere 5.5 and later, though a virtual switch (standard switch or distributed switch) does not handle the network traffic of an SR-IOV enabled virtual machine connected to the switch, you can control the assigned virtual functions by using switch configuration policies at port group or port level.

SR-IOV Support

vSphere 5.1 and later supports SR-IOV in an environment with specific configuration only. Some features of vSphere are not functional when SR-IOV is enabled.

Supported Configurations

To use SR-IOV in vSphere 5.5, your environment must meet several configuration requirements.

Table 7-1. Supported Configurations for Using SR-IOV

Component	Requirements
vSphere	■ Hosts with Intel processors require ESXi 5.1 or later. ■ Hosts with AMD processors are supported with SR-IOV in ESXi 5.5 or later.
Physical host	■ Must be compatible with the ESXi release. ■ Must have an Intel processor if you are running ESXi 5.1, or an Intel or AMD processor if you are running ESXi 5.5 and later. ■ Must support I/O memory management unit (IOMMU), and must have IOMMU enabled in the BIOS. ■ Must support SR-IOV, and must have SR-IOV enabled in the BIOS. Contact the server vendor to determine whether the host supports SR-IOV.

Table 7-1. Supported Configurations for Using SR-IOV (Continued)

Component	Requirements
Physical NIC	■ Must be compatible with the ESXi release. ■ Must be supported for use with the host and SR-IOV according to the technical documentation from the server vendor. ■ Must have SR-IOV enabled in the firmware.
PF driver in ESXi for the physical NIC	■ Must be certified by VMware. ■ Must be installed on the ESXi host. The ESXi release provides a default driver for certain NICs, while for others you must download and manually install it.
Guest OS	■ Red Hat Enterprise Linux 6.x ■ Windows Server 2008 R2 with SP2
VF driver in the guest OS	■ Must be compatible with the NIC. ■ Must be supported on the guest OS release according to the technical documentation from the NIC vendor. ■ Must be Microsoft WLK or WHCK certified for Windows virtual machines. ■ Must be installed on the operating system. The operating system release contains a default driver for certain NICs, while for others you must download and install it from a location provided by the vendor of the NIC or the host.

To verify that your physical hosts and NICs are compatible with ESXi releases, see the *VMware Compatibility Guide*.

Availability of Features

The following features are not available for virtual machines configured with SR-IOV:

- vSphere vMotion
- Storage vMotion
- vShield
- NetFlow
- VXLAN Virtual Wire
- vSphere High Availability
- vSphere Fault Tolerance
- vSphere DRS
- vSphere DPM
- Virtual machine suspend and resume
- Virtual machine snapshots
- MAC-based VLAN for passthrough virtual functions
- Hot addition and removal of virtual devices, memory, and vCPU
- Participation in a cluster environment

- Network statistics for a virtual machine NIC using SR-IOV passthrough

NOTE Attempts to enable or configure unsupported features with SR-IOV in the vSphere Web Client result in unexpected behavior in your environment.

Supported NICs

All NICs must have drivers and firmware that support SR-IOV. Some NICs might require SR-IOV to be enabled on the firmware. The following NICs are supported for virtual machines configured with SR-IOV:

- Products based on the Intel 82599ES 10 Gigabit Ethernet Controller Family (Niantic)
- Products based on the Intel Ethernet Controller X540 Family (Twinville)
- Emulex OneConnect (BE3)

Upgrading from vSphere 5.0 and Earlier

If you upgrade from vSphere 5.0 or earlier to vSphere 5.5 or later, SR-IOV support is not available until you update the NIC drivers for the vSphere release. Firmware and drivers that support SR-IOV must be enabled for NICs so that SR-IOV functionality can operate.

Upgrading from vSphere 5.1

Although SR-IOV is supported on ESXi 5.1 hosts satisfying the requirements, you cannot configure SR-IOV on them by using the vSphere Web Client. Use the `max_vfs` parameter of the NIC driver module to enable SR-IOV on these hosts. See "Enabling SR-IOV by Using Host Profiles in the vSphere Web Client or Through an ESXCLI Command," on page 140.

You also cannot assign an SR-IOV passthrough adapter to a virtual machine on such a host. The adapter is available for virtual machines that are compatible with ESXi 5.5 and later. Although a vCenter Server 5.5 release might be managing an ESXi 5.1 host, the configuration is the same as in release 5.1. You must add a PCI device to the virtual machine hardware and manually select a VF for the device.

SR-IOV Component Architecture and Interaction

vSphere SR-IOV support relies on the interaction between the virtual functions (VFs) and the physical function (PF) of the NIC port for better performance, and interaction between the driver of the PF and the host switch for traffic control.

In a host that runs virtual machine traffic on top of SR-IOV physical adapters, virtual machine adapters directly contact the virtual functions to communicate data. However, the ability to configure networks is based on the active policies for the port holding the virtual machines.

On an ESXi host without SR-IOV, the virtual switch sends external network traffic through its ports on the host from or to the physical adapter for the relevant port group. The virtual switch also applies the networking policies on managed packets.

Figure 7-1. Data and Configuration Paths in the SR-IOV Support of vSphere

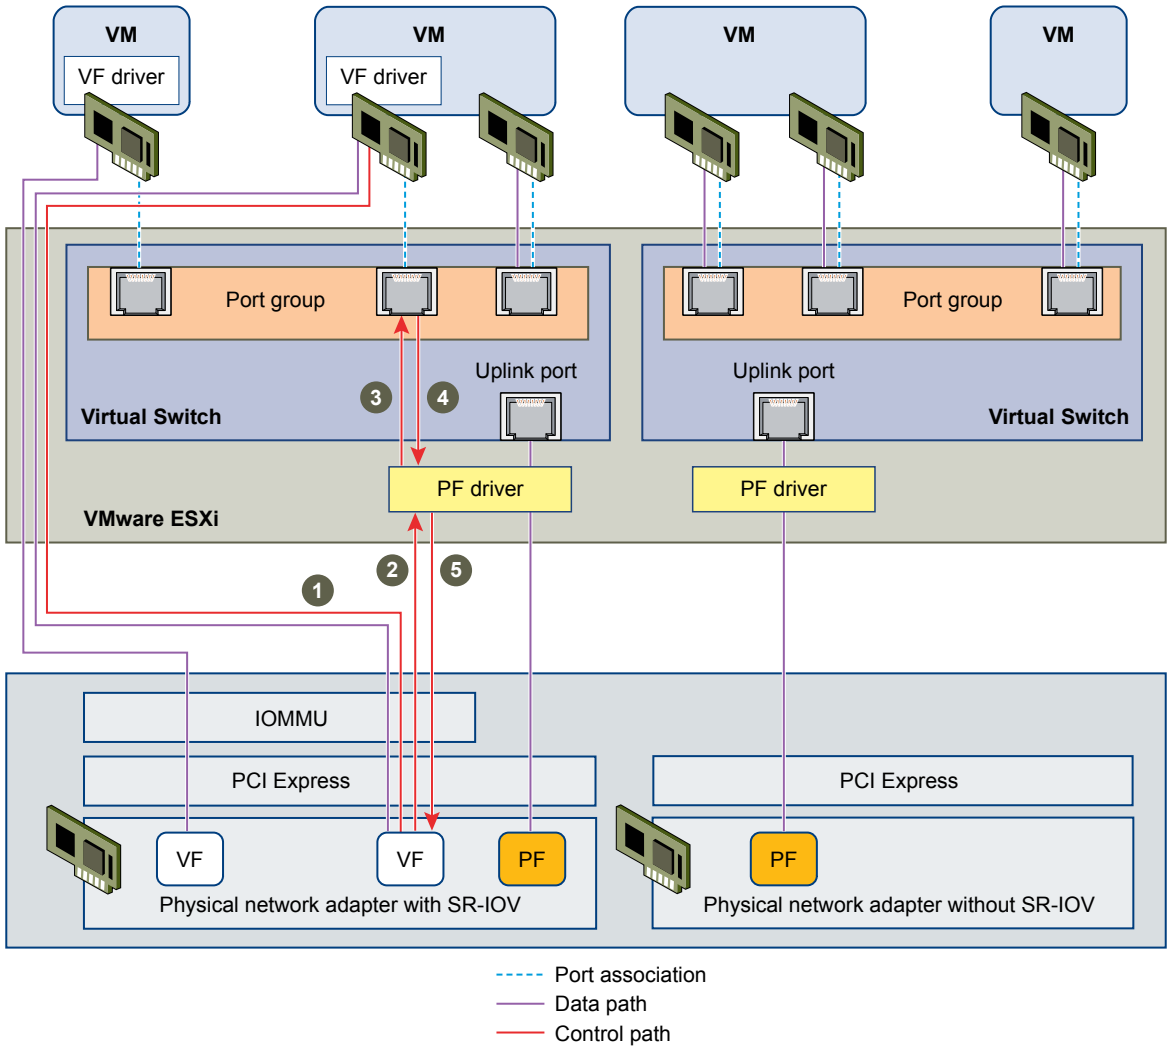

Data Path in SR-IOV

After the virtual machine network adapter is assigned to a virtual function, the VF driver in the guest operating system uses the I/O memory management unit (IOMMU) technology to access the virtual function that must receive or send the data over the network. The VMkernel, that is, the virtual switch in particular, does not process the data flow, which reduces the overall latency of SR-IOV enabled workloads.

Configuration Path in SR-IOV

If the guest operating system attempts to change the configuration of a virtual machine adapter mapped to a VF, the change takes place if it is allowed by the policy on the port associated with the virtual machine adapter.

The configuration workflow consists of the following operations:

1 The guest operating system requests a configuration change on the VF.

2 The VF forwards the request to the PF through a mailbox mechanism.

3 The PF driver checks the configuration request with the virtual switch (standard switch or host proxy switch of a distributed switch).

4 The virtual switch verifies the configuration request against the policy on the port with which the VF enabled virtual machine adapter is associated.

5 The PF driver configures the VF if the new settings are in compliance with the port policy of the virtual machine adapter.

   For example, if the VF driver tries to modify the MAC address, the address remains the same if MAC address change is not allowed in the security policy for the port group or port. The guest operating system might show that the change is successful but a log message indicates that the operation has failed. As a result, the guest operating system and the virtual device save different MAC addresses. The network interface in the guest operating system might not be able to acquire an IP address and communicate. In this case, you have to reset the interface in the guest operating system to get the latest MAC address from the virtual device and acquire an IP address.

vSphere and Virtual Function Interaction

Virtual functions (VFs) are lightweight PCIe functions that contain all the resources necessary for data exchange but have a minimized set of configuration resources. Interaction between vSphere and VFs is limited.

- VFs do not implement rate control in vSphere. Every VF can potentially use the entire bandwidth of a physical link.
- When a VF device is configured as a passthrough device on a virtual machine, the standby and hibernate functions for the virtual machine are not supported.
- The maximum number of VFs that you can create and the maximum number of VFs that you can use for passthrough are different. The maximum number of VFs that you can instantiate depends on the NIC capability and on the hardware configuration of the host. However, due to the limited number of interrupt vectors available for passthrough devices, only a limited number of all instantiated VFs can be used on an ESXi host.

  Each ESXi host has a total of 256 interrupt vectors. When the host boots, devices on the host such as storage controllers, physical network adapters, and USB controllers consume a subset of the 256 vectors. If these devices require more than 128 vectors, the maximum number of potentially supported VFs is reduced.

  For example, out of 64 VFs that can be instantiated on an Intel NIC, the host can use up to 43 VFs for passthrough (`128 / 3 = 42.6`) if all 128 interrupt vectors are available.
- If you have Intel and Emulex NICs present with SR-IOV enabled, the number of VFs available for the Intel NICs depends on how many VFs are configured for the Emulex NIC, and the reverse. You can use the following formula to estimate the maximum number of VFs for use if all 128 interrupt vectors are available for passthrough:

  ```
  3X + 2Y < 128
  ```

  where `X` is the number of Intel VFs, and `Y` is the number of Emulex VFs.

  This number might be smaller if other types of devices on the host use more than 128 interrupt vectors from the total of 256 vectors on the host.
- vSphere SR-IOV supports up to 43 VFs on supported Intel NICs and up to 64 VFs on supported Emulex NICs.
- If a supported Intel NIC loses connection, all VFs from the physical NIC stop communication completely, including that between VFs.
- If a supported Emulex NIC loses connection, all VFs stop communication with the external environment, but communication between VFs still works
- VF drivers offer many different features, such as IPv6 support, TSO, and LRO checksum. See the technical documentation of the NIC vendor for more details.

DirectPath I/O vs SR-IOV

SR-IOV offers performance benefits and tradeoffs similar to those of DirectPath I/O. DirectPath I/O and SR-IOV have similar functionality but you use them to accomplish different things.

SR-IOV is beneficial in workloads with very high packet rates or very low latency requirements. Like DirectPath I/O, SR-IOV is not compatible with certain core virtualization features, such as vMotion. SR-IOV does, however, allow for a single physical device to be shared amongst multiple guests.

With DirectPath I/O you can map only one physical function to one virtual machine. SR-IOV lets you share a single physical device, allowing multiple virtual machines to connect directly to the physical function.

Configure a Virtual Machine to Use SR-IOV in the vSphere Web Client

To use the capabilities of SR-IOV, you must enable the SR-IOV virtual functions on the host and connect a virtual machine to the functions.

Prerequisites

Verify that the configuration of your environment supports SR-IOV. See "SR-IOV Support," on page 132.

Procedure

1 Enable SR-IOV on a Host Physical Adapter in the vSphere Web Client on page 137

Before you can connect virtual machines to virtual functions, use the vSphere Web Client to enable SR-IOV and set the number of virtual functions on your host.

2 Assign a Virtual Function as SR-IOV Passthrough Adapter to a Virtual Machine by Using the vSphere Web Client on page 138

To ensure that a virtual machine and a physical NIC can exchange data, you must associate the virtual machine with one or more virtual functions as SR-IOV passthrough network adapters.

The traffic passes from an SR-IOV passthrough adapter to the physical adapter in compliance with the active policy on the associated port on the standard or distributed switch.

To examine which virtual function is assigned to an SR-IOV passthrough network adapter, on the **Summary** tab for the virtual machine expand the **VM Hardware** panel and check the properties of the adapter.

The topology diagram of the switch marks virtual machine adapters that use virtual functions with the  icon.

What to do next

Set up the traffic passing through the virtual functions attached to the virtual machine by using the networking policies on the switch, port group, and port. See "Networking Options for the Traffic Related to an SR-IOV Enabled Virtual Machine," on page 139.

Enable SR-IOV on a Host Physical Adapter in the vSphere Web Client

Before you can connect virtual machines to virtual functions, use the vSphere Web Client to enable SR-IOV and set the number of virtual functions on your host.

Procedure

1 In the vSphere Web Client, navigate to the host.

2 On the **Manage** tab, click **Networking** and select **Physical adapters**.

You can look at the SR-IOV property to see whether a physical adapter supports SR-IOV.

3 Select the physical adapter and click **Edit adapter settings**.

4 Under SR-IOV, select **Enabled** from the **Status** drop-down menu.

5 In the **Number of virtual functions** text box, type the number of virtual functions that you want to configure for the adapter.

6 Click **OK**.

7 Restart the host.

The virtual functions become active on the NIC port represented by the physical adapter entry. They appear in the PCI Devices list in the **Settings** tab for the host.

You can use the `esxcli network sriovnic` vCLI commands to examine the configuration of virtual functions on the host.

What to do next

Associate a virtual machine with a virtual function through an SR-IOV passthrough network adapter.

Assign a Virtual Function as SR-IOV Passthrough Adapter to a Virtual Machine by Using the vSphere Web Client

To ensure that a virtual machine and a physical NIC can exchange data, you must associate the virtual machine with one or more virtual functions as SR-IOV passthrough network adapters.

Prerequisites

- Verify that the virtual functions exist on the host.
- Verify that the passthrough networking devices for the virtual functions are active in the PCI Devices list on the **Settings** tab for the host.
- Verify that the virtual machine compatibility is ESXi 5.5 and later.
- Verify that Red Hat Enterprise Linux 6 and later or Windows has been selected as the guest operating system when the virtual machine was created.

Procedure

1 Locate the virtual machine in the vSphere Web Client.

   a Select a datacenter, folder, cluster, resource pool, or host and click the **Related Objects** tab.

   b Click **Virtual Machines** and select the virtual machine from the list.

2 Power off the virtual machine.

3 On the **Manage** tab of the virtual machine, select **Settings > VM Hardware**.

4 Click **Edit** and select the **Virtual Hardware** tab.

5 From the **New device** drop-down menu, select **Network** and click **Add**.

6 Expand the New Network section and connect the virtual machine to a port group.

   The virtual NIC does not use this port group for data traffic. The port group is used to extract the networking properties, for example VLAN tagging, to apply on the data traffic.

7 From the **Adapter type** drop-down menu, select **SR-IOV passthrough**.

8 From the **Physical function** drop-down menu, select the physical adapter to back the passthrough virtual machine adapter.

9 To allow changes in the MTU of packets from the guest operating system, use the **Guest OS MTU Change** drop-down menu.

10 Expand the Memory section, select **Reserve all guest memory (All locked)** and click **OK**.

I/O memory management unit (IOMMU) must reach all virtual machine memory so that the passthrough device can access the memory by using direct memory access (DMA).

11 Power on the virtual machine.

When you power on the virtual machine, the ESXi host selects a free virtual function from the physical adapter and maps it to the SR-IOV passthrough adapter. The host validates all properties of the virtual machine adapter and the underlying virtual function against the settings of the port group to which the virtual machine belongs.

Networking Options for the Traffic Related to an SR-IOV Enabled Virtual Machine

In vSphere 5.5 and later, you can configure certain networking features on a virtual machine adapter that is associated virtual function (VF). Use settings for the switch, for the port group, or for a port depending on the type of the virtual switch (standard or distributed) that handles the traffic.

Table 7-2. Networking Options for a Virtual Machine Adapter That Uses a VF

Networking Option	Description
MTU size	Change the size of the MTU, for example, to enable jumbo frames.
Security policy for VF traffic	■ If the guest operating system changes the initially set MAC address of a virtual machine network adapter that uses a VF, accept or drop incoming frames for the new address by setting the **MAC address changes** option. ■ Enable global promiscuous mode for virtual machine network adapters, including adapters that use VFs.
VLAN tagging mode	Configure VLAN tagging in the standard or distributed switch, that is, enable VLAN Switch Tagging (VST) mode, or let the tagged traffic reach the virtual machines that are associated with VFs, that is, enable Virtual Guest Tagging (VGT).

Using an SR-IOV Physical Adapter to Handle Virtual Machine Traffic

In vSphere 5.5 and later, both the physical function (PF) and virtual functions (VFs) of an SR-IOV capable physical adapter can be configured to handle virtual machine traffic.

The PF of an SR-IOV physical adapter controls the VFs that virtual machines use, and can carry the traffic flowing through the standard or distributed switch that handles the networking of these SR-IOV enabled virtual machines.

The SR-IOV physical adapter works in different modes depending on whether it backs the traffic of the switch.

Mixed Mode

The physical adapter provides virtual functions to virtual machines attached to the switch and directly handles traffic from non SR-IOV virtual machines on the switch.

You can check whether an SR-IOV physical adapter is in mixed mode in the topology diagram of the switch.

An SR-IOV physical adapter in mixed mode appears with the  icon in the list of physical adapters for a standard switch or in the list of uplink group adapters for a distributed switch.

SR-IOV Only Mode

The physical adapter provides virtual functions to virtual machines connected to a virtual switch, but does not back traffic from non SR-IOV virtual machines on the switch.

To verify whether the physical adapter is in SR-IOV only mode, examine the topology diagram of the switch. In this mode, the physical adapter is in a separate list called External SR-IOV Adapters and appears with the  icon.

Non SR-IOV Mode

The physical adapter is not used for traffic related to VF aware virtual machines. It handles traffic from non SR-IOV virtual machines only.

Enabling SR-IOV by Using Host Profiles in the vSphere Web Client or Through an ESXCLI Command

You can configure the virtual functions on an ESXi host by using an ESXCLI command, or by using a host profile to set up multiple hosts simultaneously or to set up stateless hosts.

Enable SR-IOV in a Host Profile by Using the vSphere Web Client

For multiple hosts or a stateless host, you can configure the virtual functions of the physical NIC by using a host profile and apply the profile on a host by using Auto Deploy.

For information about running ESXi by using Auto Deploy with host profiles, see the *vSphere Installation and Setup* documentation.

You can also enable SR-IOV virtual functions on the host by using the `esxcli system module parameters set` vCLI command on the NIC driver parameter for virtual functions in accordance with the driver documentation. For more information about using vCLI commands, see *vSphere Command-Line Interface Documentation*.

Prerequisites

- Verify that the configuration of your environment supports SR-IOV. See "SR-IOV Support," on page 132.
- Create a host profile based on the SR-IOV capable host. See the *vSphere Host Profiles* documentation.

Procedure

1. From the vSphere Web Client Home page, click **Rules and Profiles > Host Profiles**.
2. Select the host profile from the list and click the **Manage** tab.
3. Click **Edit Host Profile** and expand the **General System Settings** node.
4. Expand **Kernel Module Parameter** and select the parameter of the physical function driver for creating virtual functions.

   For example, the parameter for the physical function driver of an Intel physical NIC is `max_vfs`.
5. In the **Value** text box, type a comma-separated list of valid virtual function numbers.

   Each list entry indicates the number of virtual functions that you want to configure for each physical function. A value of 0 ensures that SR-IOV is not enabled for that physical function.

   For example, if you have a dual port, set the value to *x*,*y* where *x* or *y* is the number of virtual functions you want to enable for a single port.

If the target number of virtual functions on a single host is 30, you might have two dual port cards set to `0,10,10,10`.

NOTE The number of virtual functions supported and available for configuration depends on your system configuration.

6 Click **Finish**.

7 Remediate the host profile to the host as required.

The virtual functions appear in the PCI Devices list on the **Settings** tab for the host.

What to do next

Associate a virtual function with a virtual machine adapter by using the SR-IOV passthrough network adapter type. See "Assign a Virtual Function as SR-IOV Passthrough Adapter to a Virtual Machine by Using the vSphere Web Client," on page 138.

Enable SR-IOV on a Host Physical Adapter by Using an ESXCLI Command

In certain troubleshooting situations or to configure hosts directly, you can run a console command on ESXi to create SR-IOV virtual functions on a physical adapter.

You can create SR-IOV virtual functions on the host by manipulating the NIC driver parameter for virtual functions in accordance with the driver documentation.

Prerequisites

Install the vCLI package, deploy the vSphere Management Assistant (vMA) virtual machine, or use the ESXi Shell. See *Getting Started with vSphere Command-Line Interfaces*.

Procedure

1 To create virtual functions by setting the parameter for virtual functions of the NIC driver, run the `esxcli system module parameters set` command at the command prompt.

```
esxcli system    module parameters set –m driver –p vf_param=w,x,y,z
```

Where *driver* is the name of the NIC driver, and *vf_param* is the driver-specific parameter for creating the virtual function.

You can use a comma-separated list to set values for the *vf_param* parameter, where each entry indicates the number of virtual functions for a port. A value of 0 ensures that SR-IOV is not enabled for that physical function.

If you have two dual port NICs, you can set the value to *w*,*x*,*y*,*z*, where *w*,*x*,*y*, and *z* is the number of virtual functions you want to enable for a single port. For example, to create 30 virtual functions distributed on two dual port Intel cards by using the ixgbe driver, run the following command for the ixgbe driver and the `max_vfs` parameter:

```
esxcli system module parameters set –m ixgbe –p max_vfs=0,10,10,10
```

2 Restart the host to create the virtual functions.

What to do next

Associate a virtual function with a virtual machine adapter by using the SR-IOV passthrough network adapter type. See "Assign a Virtual Function as SR-IOV Passthrough Adapter to a Virtual Machine by Using the vSphere Web Client," on page 138.

A Virtual Machine That Uses an SR-IOV Virtual Function Is Powered off Because the Host Is Out of Interrupt Vectors

On an ESXi host, one or more virtual machines that use SR-IOV virtual functions (VFs) for networking are powered off.

Problem

On an ESXi host, one or more virtual machines that use SR-IOV virtual functions (VFs) for networking are powered off when the total number of assigned virtual functions is close to the maximum number of VFs specified in the *vSphere Configuration Maximums* guide.

The virtual machine log file `vmware.log` contains the following message about the VF:

```
PCIPassthruChangeIntrSettings: vf_name failed to register interrupt (error code 195887110)
```

The VMkernel log file `vmkernel.log` contains the following messages about the VF assigned to the virtual machine:

```
VMKPCIPassthru: 2565: BDF = vf_name intrType = 4 numVectors: 3
WARNING: IntrVector: 233: Out of interrupt vectors
```

Cause

Each ESXi host has a total of 256 interrupt vectors. When the host boots, devices on the host such as storage controllers, physical network adapters, and USB controllers consume a subset of the 256 vectors. If these devices require more than 128 vectors, the maximum number of potentially supported VFs is reduced.

When a virtual machine powers on and the guest operating system VF driver starts, interrupt vectors are consumed. If the required number of interrupt vectors is not available, the guest operating system shuts down unexpectedly without any error messages.

No method presently exists to determine the number of interrupt vectors consumed or available on a host. This number depends on the hardware configuration of the host.

Solution

To be able to power on the virtual machines, reduce the total number of VFs assigned to virtual machines on the host. For example, change the SR-IOV network adapter of a virtual machine to an adapter that is connected to a vSphere Standard Switch or vSphere Distributed Switch.

8 MAC Address Management

MAC addresses are used in the Layer 2 (Data Link Layer) of the network protocol stack to transmit frames to a recipient. In vSphere, vCenter Server generates MAC addresses for virtual machine adapters and VMkernel adapters, or you can assign addresses manually.

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 supports several address allocation mechanisms, each of them with a separate OUI:

- Generated MAC addresses
  - Assigned by vCenter Server
  - Assigned by the ESXi host
- Manually set MAC addresses
- Generated for legacy virtual machines, but no longer used with ESXi

If you reconfigure the network adapter of a powered off virtual machine, for example by changing the automatic MAC address allocation type, or setting a static MAC address, vCenter Server resolves any MAC address conflict before the adapter reconfiguration takes effect.

This chapter includes the following topics:

MAC Address Assignment from vCenter Server

vSphere 5.1 and later provides several schemes for automatic allocation of MAC addresses in vCenter Server. You can select the scheme that best suits your requirements for MAC address duplication, OUI requirements for locally administered or universally administered addresses, and so on.

The following schemes of MAC address generation are available in vCenter Server:

- VMware OUI allocation, default allocation
- Prefix-based allocation
- Range-based allocation

After the MAC address is generated, it does not change unless the virtual machine's MAC address conflicts with that of another registered virtual machine. The MAC address is saved in the configuration file of the virtual machine.

NOTE If you use invalid prefix- or range-based allocation values, an error is logged in the `vpxd.log` file. vCenter Server does not allocate MAC addresses when provisioning a virtual machine.

Preventing MAC Address Conflicts

The MAC address of a powered off virtual machine is not checked against the addresses of running or suspended virtual machines.

When a virtual machine is powered on again, it might acquire a different MAC address. The change might be caused by an address conflict with another virtual machine. While this virtual machine has been powered off, its MAC address has been assigned to another virtual machine that has been powered on.

If you reconfigure the network adapter of a powered off virtual machine, for example, by changing the automatic MAC address allocation type or setting a static MAC address, vCenter Server resolves MAC address conflicts before the adapter reconfiguration takes effect.

For information about resolving MAC address conflicts, see the *vSphere Troubleshooting* documentation.

VMware OUI Allocation

VMware Organizationally Unique Identifier (OUI) allocation assigns MAC addresses based on the default VMware OUI `00:50:56` and the vCenter Server ID.

VMware OUI allocation is the default MAC address assignment model for virtual machines. The allocation works with up to 64 vCenter Server instances, and each vCenter Server can assign up to 64000 unique MAC addresses. The VMware OUI allocation scheme is suitable for small scale deployments.

MAC Address Format

According to the VMware OUI allocation scheme, a MAC address has the format `00:50:56:`*`XX`*`:`*`YY`*`:`*`ZZ`* where `00:50:56` represents the VMware OUI, *XX* is calculated as (80 + vCenter Server ID), and *YY* and *ZZ* are random two-digit hexadecimal numbers.

The addresses created through the VMware OUI allocation are in the range `00:50:56:80:`*`YY`*`:`*`ZZ`* - `00:50:56:BF:`*`YY`*`:`*`ZZ`*.

Prefix-Based MAC Address Allocation

On ESXi hosts 5.1 and later, you can use prefix-based allocation to specify an OUI other than the default one `00:50:56` by VMware, or to introduce Locally Administered MAC Addresses (LAA) for a larger address space.

Prefix-based MAC address allocation overcomes the limits of the default VMware allocation to provide unique addresses in larger scale deployments. Introducing an LAA prefix leads to a very large MAC address space (2 to the power of 46) instead of an universally unique address OUI which can give only 16 million MAC addresses.

Verify that the prefixes that you provide for different vCenter Server instances in the same network are unique. vCenter Server relies on the prefixes to avoid MAC address duplication issues. See the *vSphere Troubleshooting* documentation.

Range-Based MAC Address Allocation

On ESXi hosts 5.1 and later you can use range-based allocation to include or exclude ranges of Locally Administered Addresses (LAA).

You specify one or more ranges using a starting and ending MAC addresses, for example, (`02:50:68:00:00:02`, `02:50:68:00:00:FF`). MAC addresses are generated only from within the specified range.

You can specify multiple ranges of LAA, and vCenter Server tracks the number of used addresses for each range. vCenter Server allocates MAC addresses from the first range that still has addresses available. vCenter Server checks for MAC address conflicts within its ranges.

When using range-based allocation, you must provide different instances of vCenter Server with ranges that do not overlap. vCenter Server does not detect ranges that might be in conflict with other vCenter Server instances. See the *vSphere Troubleshooting* documentation for more information about resolving issues with duplicate MAC addresses.

Assigning a MAC Address

Use the vSphere Web Client to enable prefix-based or range-based MAC address allocation and to adjust the allocation parameters.

If you are changing from one type of allocation to another, for example changing from the VMware OUI allocation to a range-based allocation, use the vSphere Web Client. However, when a schema is prefix-based or range-based and you want to change to a different allocation schema, you must edit the `vpxd.cfd` file manually and restart vCenter Server.

Change to or Adjust Range- or Prefixed-Based Allocations in the vSphere Web Client

By switching from the default VMware OUI to range- or prefixed-based MAC address allocation through the vSphere Web Client, you can avoid and resolve MAC address duplication conflicts in vSphere deployments.

Change the allocation scheme from the default VMware OUI to range- or to prefixed-based allocation by using the **Advanced Settings** available for the vCenter Server instance in the vSphere Web Client.

To switch from range- or prefixed-based allocation back to VMware OUI allocation, or between range- and prefixed-based allocation, edit the `vpxd.cfg` file manually. See "Set or Change Allocation Type," on page 146.

NOTE You should use prefix-based MAC address allocation in vCenter Server 5.1 and ESXi 5.1 hosts, and later.

If a vCenter Server 5.1 instance manages hosts running ESXi versions earlier than ESXi 5.1, use VMware OUI prefix-based MAC address allocation. Virtual machines assigned non VMware OUI prefixed MAC addresses fail to power on on the pre-5.1 hosts. These hosts explicitly check if an assigned MAC address uses the VMware OUI 00:50:56 prefix.

Procedure

1 Browse to a vCenter Server in the vSphere Web Client.

2 Click the **Manage** tab, and select **Settings > Advanced Settings**.

3 Click **Edit.**

4 Add or edit parameters for the target allocation type.

Use only one allocation type.

- Change to prefix-based allocation.

Key	Example Value
`config.vpxd.macAllocScheme.prefixScheme.prefix`	005026
`config.vpxd.macAllocScheme.prefixScheme.prefixLength`	23

`prefix` and `prefixLength` determine the range of MAC address prefixes that newly added vNICs have. `prefix` is the starting OUI of MAC addresses related to the vCenter Server instance, and `prefixLength` determines the length of the prefix in bits.

For example, the settings from the table result in vNIC MAC addresses starting with 00:50:26 or 00:50:27.

- Change to range-based allocation.

Key	Example Value
`config.vpxd.macAllocScheme.rangeScheme.range[X].begin`	005067000000
`config.vpxd.macAllocScheme.rangeScheme.range[X].end`	005067ffffff

X in `range[X]` stands for the range sequence number. For example, `0` in `range[0]` represents the allocation settings of the first range for MAC address allocation.

5 Click **OK**.

Set or Change Allocation Type

If you are changing from range- or prefixed-based allocation to the VMware OUI allocation, you must set the allocation type in the `vpxd.cfd` file and restart the vCenter Server.

Prerequisites

Decide on an allocation type before changing the `vpxd.cfg` file. For information on allocation types, see "MAC Address Assignment from vCenter Server," on page 143

Procedure

1 On the host machine of vCenter Server, navigate to the directory that contains the configuration file:

- On a Windows Server operating system, the location of the directory is *vCenter Server home directory*`\Application Data\VMware\VMware VirtualCenter`.
- On the vCenter Server Appliance, the location of the directory is `/etc/vmware-vpx`.

2 Open the `vpxd.cfg` file.

3 Decide on an allocation type to use and enter the corresponding XML code in the file to configure the allocation type.

The following are examples of XML code to use.

NOTE Use only one allocation type.

- VMware OUI allocation

```
<vpxd>
   <macAllocScheme>
      <VMwareOUI>true</VMwareOUI>
   </macAllocScheme>
</vpxd>
```

- Prefix-based allocation

```
 <vpxd>
    <macAllocScheme>
      <prefixScheme>
         <prefix>005026</prefix>
         <prefixLength>23</prefixLength>
      </prefixScheme>
   </macAllocScheme>
</vpxd>
```

- Range-based allocation

```
<vpxd>
   <macAllocScheme>
       <rangeScheme>
          <range id="0">
             <begin>005067000001</begin>
             <end>005067000001</end>
          </range>
       </rangeScheme>
   </macAllocScheme>
</vpxd>
```

4 Save the `vpxd.cfg`.

5 Restart the vCenter Server host.

MAC Address Generation on ESXi Hosts

An ESXi host generates the MAC address for a virtual machine adapter when the host is not connected to vCenter Server. Such addresses have a separate VMware OUI to avoid conflicts.

The ESXi host generates the MAC address for a virtual machine adapter in one of the following cases:

- The host is not connected to vCenter Server.
- The virtual machine configuration file does not contain the MAC address and information about the MAC address allocation type.

MAC Address Format

The host generates MAC addresses that consists of the VMware OUI `00:0C:29` and the last three octets in hexadecimal format of the virtual machine UUID. The virtual machine UUID is based on a hash calculated by using the UUID of the ESXi physical machine and the path to the configuration file (`.vmx`) of the virtual machine.

Preventing MAC Address Conflicts

All MAC addresses that have been assigned to network adapters of running and suspended virtual machines on a given physical machine are tracked for conflicts.

If you import a virtual machine with a host-generated MAC address from one vCenter Server to another, select the **I Copied It** option when you power on the virtual machine to regenerate the address and avoid potential conflicts in the target vCenter Server or between the vCenter Server systems.

Setting a Static MAC Address to a Virtual Machine

In most network deployments, generated MAC addresses are a good approach. However, you might need to set a static MAC address for a virtual machine adapter with unique value.

The following cases show when you might set a static MAC address:

- Virtual machine adapters on different physical hosts share the same subnet and are assigned the same MAC address, causing a conflict.
- Ensure that a virtual machine adapter always has the same MAC address.

By default, VMware uses the Organizationally Unique Identifier (OUI) `00:50:56` for manually generated addresses, but all unique manually generated addresses are supported.

NOTE Make sure that no other non-VMware devices use addresses assigned to VMware components. For example, you might have physical servers in the same subnet, which use `11:11:11:11:11:11`, `22:22:22:22:22:22` as static MAC addresses. The physical servers do not belong to the vCenter Server inventory, and vCenter Server is not able to check for address collision.

VMware OUI in Static MAC Addresses

By default, static MAC addresses have the VMware Organizationally Unique Identifier (OUI) as the prefix. However, the range of free address provided by the VMware OUI is restrictions.

If you choose to use the VMware OUI, part of the range is reserved for use by vCenter Server, host physical NICs, virtual NICs, and for future use.

You can set a static MAC address that contains the VMware OUI prefix in compliance with the following format:

`00:50:56:`*XX*`:`*YY*`:`*ZZ*

where *XX* is a valid hexadecimal number between 00 and 3F, and *YY* and *ZZ* are valid hexadecimal numbers between 00 and FF. To avoid conflict with MAC addresses that are generated by vCenter Server or are assigned to VMkernel adapters for infrastructure traffic, the value for *XX* must not be greater than 3F.

The maximum value for a manually generated MAC address is as follows.

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

To avoid conflicts between the generated MAC addresses and the manually assigned ones, select a unique value for *XX*:*YY*:*ZZ* from your hard-coded addresses.

Assign a Static MAC Address with the vSphere Web Client

You can assign static MAC addresses to the virtual NIC of a powered off virtual machine by using the vSphere Web Client.

Procedure

1 Locate the virtual machine in the vSphere Web Client.

   a Select a datacenter, folder, cluster, resource pool, or host and click the **Related Objects** tab.

   b Click **Virtual Machines** and select the virtual machine from the list.

2 Power off the virtual machine.

3 On the **Manage** tab of the virtual machine, select **Settings > VM Hardware**.

4 Click **Edit** and select the **Virtual Hardware** tab.

5 In the **Virtual Hardware** tab, expand the network adapter section.

6 Under MAC Address, select **Manual** from the drop-down menu.

7 Type the static MAC address, and click **OK**.

8 Power on the virtual machine.

Assign a Static MAC Address in the Virtual Machine Configuration File

To set a static MAC address for a virtual machine, you can edit the configuration file of the virtual machine by using the vSphere Web Client.

Procedure

1 Locate the virtual machine in the vSphere Web Client.

   a Select a datacenter, folder, cluster, resource pool, or host and click the **Related Objects** tab.

   b Click **Virtual Machines** and select the virtual machine from the list.

2 Power off the virtual machine.

3 On the **Manage** tab of the virtual machine, select **Settings**.

4 On the **VM Options** tab, expand **Advanced**.

5 Click **Edit Configuration**.

6 To assign a static MAC address, add or edit parameters as required.

Parameter	Value
`ethernet`*X*`.addressType`	static
`ethernet`*X*`.address`	*MAC_address_of_the_virtual_NIC*

   X next to `ethernet` stands for the sequence number of the virtual NIC in the virtual machine.

   For example, `0` in `ethernet0` represents the settings of the first virtual NIC device added to the virtual machine.

7 Click **OK**.

8 Power on the virtual machine.

Advanced Networking 9

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

This chapter includes the following topics:

Enable or Disable IPv6 Support on a Host by Using the vSphere Web Client

The IPv6 support in vSphere lets hosts work in an IPv6 network that has a very large address space, enhanced multicasting, simplified routing, and so on.

IPv6 is designated by the Internet Engineering Task Force as the successor to IPv4. The most obvious advantage of IPv6 to IPv4 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.

In ESXi 5.1 and later releases, IPv6 is enabled by default.

Prerequisites

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

Procedure

1 In the vSphere Web Client, navigate to the host.

2 On the **Manage** tab, click **Networking** and select **Advanced**.

3 Click **Edit**.

4 From the **IPv6 support** drop-down menu, enable or disable IPv6 support.

5 Click **OK**.

6 Reboot the host to apply the changes in the IPv6 support.

What to do next

Configure the IPv6 settings of VMkernel adapters, for example, of the management network. See "Edit a VMkernel Adapter Configuration in the vSphere Web Client," on page 72.

Working With Port Mirroring

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

Port mirroring is used on a switch to send a copy of packets seen on one switch port (or an entire VLAN) to a monitoring connection on another switch port. Port mirroring is used to analyze and debug data or diagnose errors on a network.

Port Mirroring Version Compatibility

Certain port mirroring functionality in vSphere 5.1 and later depends on which version of vCenter Server, vSphere distributed switch, and host you use, and how you use these aspects of vSphere together.

Table 9-1. Port mirroring compatibility

vCenter Server version	vSphere distributed switch version	Host version	vSphere 5.1 port mirroring functionality
vSphere 5.1 and later	vSphere 5.1 and later	vSphere 5.1 and later	vSphere 5.1 port mirroring is available for use. Features for vSphere 5.0 and earlier port mirroring are not available.
vSphere 5.1 and later	vSphere 5.1 and later	vSphere 5.0 and earlier	vSphere 5.0 and earlier hosts can be added to vSphere 5.1 vCenter Server, but cannot be added to distributed switches version 5.1 and later.
vSphere 5.1 and later	vSphere 5.0	vSphere 5.0	vSpherevCenter Server version 5.1 and later can configure port mirroring on a vSphere 5.0 distributed switch.
vSphere 5.1 and later	vSphere 5.0	vSphere 5.1 and later	Hosts running vSphere 5.1 can be added to vSphere 5.0 distributed switches and support vSphere 5.0 port mirroring.
vSphere 5.1 and later	Pre-vSphere 5.0	vSphere 5.5 and earlier	Port mirroring is not supported.
vSphere 5.0 and earlier	vSphere 5.0 and earlier	vSphere 5.1	A vSphere 5.1 host cannot be added to vCenter Server 5.0 and earlier.

If you use a host profile with port mirroring settings, the host profile must be adapted to the new version of port mirroring in vSphere 5.1 and later.

Port Mirroring Interoperability

There are some interoperability issues to consider when using vSphere 5.1 port mirroring with other features of vSphere.

vMotion

vMotion functions differently depending on which vSphere 5.1 port mirroring session type you select. During vMotion, a mirroring path could be temporarily invalid, but is restored when vMotion completes.

Table 9-2. vMotion Interoperability with port mirroring

Port mirroring session type	Source and destination	Interoperable with vMotion	Functionality
Distributed Port Mirroring	Non-uplink distributed port source and destination	Yes	Port mirroring between distributed ports can only be local. If the source and destination are on different hosts due to vMotion, mirroring between them will not work. However, if the source and destination move to the same host, port mirroring works.
Remote Mirroring Source	Non-uplink distributed port source	Yes	When a source distributed port is moved from host A to host B, the original mirroring path from the source port to A's uplink is removed on A, and a new mirroring path from the source port to B's uplink is created on B. Which uplink is used is determined by the uplink name specified in session.
	Uplink port destinations	No	Uplinks can not be moved by vMotion.
Remote Mirroring Destination	VLAN source	No	
	Non-uplink distributed port destination	Yes	When a destination distributed port is moved from host A to host B, all original mirroring paths from source VLANs to the destination port are moved from A to B.
Encapsulated Remote Mirroring (L3) Source	Non-uplink distributed port source	Yes	When a source distributed port is moved from host A to host B, all original mirroring paths from the source port to destination IPs are moved from A to B.
	IP destination	No	

Table 9-2. vMotion Interoperability with port mirroring (Continued)

Port mirroring session type	Source and destination	Interoperable with vMotion	Functionality
Distributed Port Mirroring (legacy)	IP source	No	
	Non-uplink distributed port destination	No	When a destination distributed port is moved from host A to host B, all original mirroring paths from source IPs to the destination port are invalid because the port mirroring session source still sees the destination on A.

TSO and LRO

TCP Segmentation Offload (TSO) and large receive offload (LRO) might cause the number of mirroring packets to not equal to the number of mirrored packets.

When TSO is enabled on a vNIC, the vNIC might send a large packet to a distributed switch. When LRO is enabled on a vNIC, small packets sent to it might be merged into a large packet.

Source	Destination	Description
TSO	LRO	Packets from the source vNIC might be large packets, and whether they are split is determined by whether their sizes are larger than the destination vNIC LRO limitation.
TSO	Any destination	Packets from the source vNIC might be large packets, and they are split to standard packets at the destination vNIC.
Any source	LRO	Packets from the source vNIC are standard packets, and they might be merged into larger packets at the destination vNIC.

Create a Port Mirroring Session with the vSphere Web Client

Create a port mirroring session with the vSphere Web Client to mirror vSphere distributed switch traffic to ports, uplinks, and agent's remote IP addresses.

Prerequisites

Create a vSphere distributed switch version 5.0.0 or later.

Procedure

1 Select Port Mirroring Session Type with the vSphere Web Client on page 155

To begin a port mirroring session, you must specify the type of port mirroring session.

2 Specify Port Mirroring Name and Session Details with the vSphere Web Client on page 155

To continue creating a port mirroring session, specify the name, description, and session details for the new port mirroring session.

3 Select Port Mirroring Sources with the vSphere Web Client on page 156

To continue creating a port mirroring session, select sources and traffic direction for the new port mirroring session.

4 Select Port Mirroring Destinations and Verify Settings with the vSphere Web Client on page 156

To complete the creation of a port mirroring session, select ports or uplinks as destinations for the port mirroring session.

Select Port Mirroring Session Type with the vSphere Web Client

To begin a port mirroring session, you must specify the type of port mirroring session.

Procedure

1 Browse to a distributed switch in the vSphere Web Client navigator.

2 Click the **Manage** tab and select **Settings > Port Mirroring**

3 Click **New**.

4 Select the session type for the port mirroring session.

Option	Description
Distributed Port Mirroring	Mirror packets from a number of distributed ports to other distributed ports on the same host. If the source and the destination are on different hosts, this session type does not function.
Remote Mirroring Source	Mirror packets from a number of distributed ports to specific uplink ports on the corresponding host.
Remote Mirroring Destination	Mirror packets from a number of VLANs to distributed ports.
Encapsulated Remote Mirroring (L3) Source	Mirror packets from a number of distributed ports to remote agent's IP addresses. The virtual machine's traffic is mirrored to a remote physical destination through an IP tunnel.
Distributed Port Mirroring (legacy)	Mirror packets from a number of distributed ports to a number of distributed ports and/or uplink ports on the corresponding host.

5 Click **Next**.

Specify Port Mirroring Name and Session Details with the vSphere Web Client

To continue creating a port mirroring session, specify the name, description, and session details for the new port mirroring session.

Procedure

1 Set the session properties. Different options are available for configuration depending on which session type you selected.

Option	Description
Name	You can enter a unique name for the port mirroring session, or accept the automatically generated session name.
Status	Use the drop down menu to enable or disable the session.
Session type	Displays the type of session you selected.
Normal I/O on destination ports	Use the drop-down menu to allow or disallow normal I/O on destination ports. This property is only available for uplink and distributed port destinations. If you disallow this option, mirrored traffic will be allowed out on destination ports, but no traffic will be allowed in.
Mirrored packet length (Bytes)	Use the check box to enable mirrored packet length in bytes. This puts a limit on the size of mirrored frames. If this option is selected, all mirrored frames are truncated to the specified length.
Sampling rate	Select the rate at which packets are sampled. This is enabled by default for all port mirroring sessions except legacy sessions.
Description	You have the option to enter a description of the port mirroring session configuration.

2 Click **Next**.

Select Port Mirroring Sources with the vSphere Web Client

To continue creating a port mirroring session, select sources and traffic direction for the new port mirroring session.

You can create a port mirroring session without setting the source and destinations. When a source and destination are not set, a port mirroring session is created without the mirroring path. This allows you to create a port mirroring session with the correct properties set. Once the properties are set, you can edit the port mirroring session to add the source and destination information.

Procedure

1 Select the source of the traffic to be mirrored and the traffic direction.

Depending on the type of port mirroring session you selected, different options are available for configuration.

Option	Description
Add existing ports from a list	Click **Select distributed ports**. A dialog box displays a list of existing ports. Select the check box next to the distributed port and click **OK**. You can choose more than one distributed port.
Add existing ports by port number	Click **Add distributed ports**, enter the port number and click **OK.**
Set the traffic direction	After adding ports, select the port in the list and click the ingress, egress, or ingress/egress button. Your choice appears in the Traffic Direction column.
Specify the source VLAN	If you selected a Remote Mirroring Destination sessions type, you must specify the source VLAN. Click**Add** to add a VLAN ID. Edit the ID by using the up and down arrows, or clicking in the field and entering the VLAN ID manually.

2 Click **Next**.

Select Port Mirroring Destinations and Verify Settings with the vSphere Web Client

To complete the creation of a port mirroring session, select ports or uplinks as destinations for the port mirroring session.

You can create a port mirroring session without setting the source and destinations. When a source and destination are not set, a port mirroring session is created without the mirroring path. This allows you to create a port mirroring session with the correct properties set. Once the properties are set, you can edit the port mirroring session to add the source and destination information.

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 Select the destination for the port mirroring session.

Depending on which type of session you chose, different options are available.

Option	Description
Select a destination distributed port	Click **Select distributed ports** to select ports from a list, or click **Add distributed ports** to add ports by port number. You can add more than one distributed port.
Select an uplink	Select an available uplink from the list and click **Add** to add the uplink to the port mirroring session. You can select more than one uplink.

Option	Description
Select ports or uplinks	Click **Select distributed ports** to select ports from a list, or click **Add distributed ports** to add ports by port number. You can add more than one distributed port. Click **Add uplinks** to add uplinks as the destination. Select uplinks from the list and click **OK**.
Specify IP address	Click **Add**. A new list entry is created. Select the entry and either click **Edit** to enter the IP address, or click directly in the IP Address field and type the IP address. A warning appears if the IP address is invalid.

2 Click **Next**.

3 Review the information that you entered for the port mirroring session on the **Ready to complete** page.

4 (Optional) Use the **Back** button to edit the information.

5 Click **Finish**.

The new port mirroring session appears in the Port Mirroring section of the **Settings** tab.

View Port Mirroring Session Details in the vSphere Web Client

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

Procedure

1 Browse to a distributed switch in the vSphere Web Client navigator.

2 On the **Manage** tab click **Settings > Port Mirroring**.

3 Select a port mirroring session from the list to display more detailed information at the bottom of the screen. Use the tabs to review configuration details.

4 (Optional) Click **New** to add a new port mirroring session.

5 (Optional) Click **Edit** to edit the details for the selected port mirroring session.

6 (Optional) Click **Remove** to delete the selected port mirroring session.

Edit Port Mirroring Session Details, Sources, and Destinations with the vSphere Web Client

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

Procedure

1 Browse to a distributed switch in the vSphere Web Client navigator.

2 Click the **Manage** tab and select **Settings > Port Mirroring**.

3 Select a port mirroring session from the list and click **Edit**.

4 On the **Properties** page, edit the session properties.

Depending on the type of port mirroring session being edited, different options are available for configuration.

Option	Description
Name	You can enter a unique name for the port mirroring session, or accept the automatically generated session name.
Status	Use the drop-down menu to enable or disable the session.

Option	Description
Normal I/O on destination ports	Use the drop-down menu to allow or disallow normal I/O on destination ports. This property is only available for uplink and distributed port destinations. If you do not select this option, mirrored traffic will be allowed out on destination ports, but no traffic will be allowed in.
Encapsulated VLAN ID	Enter a valid VLAN ID in the field. This information is required for Remote Mirroring Source port mirroring sessions. Mark the check box next to **Preserve original 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.
Mirrored packet length (Bytes)	Use the check box to enable mirrored packet length in bytes. This puts a limit on the size of mirrored frames. If this option is selected, all mirrored frames are truncated to the specified length.
Description	You have the option to enter a description of the port mirroring session configuration.

5 On the **Sources** page, edit sources for the port mirroring session.

Depending on the type of port mirroring session being edited, different options are available for configuration.

Option	Description
Add existing ports from a list	Click the **Select distributed ports…** button. A dialog opens with a list of existing ports. Select the check box next to the distributed port and click **OK**. You can choose more than one distributed port.
Add existing ports by port number	Click the **Add distributed ports…** button, enter the port number and click **OK.**
Set the traffic direction	After adding ports, select the port in the list and click the ingress, egress, or ingress/egress button. Your choice is displayed in the Traffic Direction column.
Specify the source VLAN	If you selected a Remote Mirroring Destination sessions type, you must specify the source VLAN. Click the **Add** button to add a VLAN ID. Edit the ID by either using the up and down arrows, or clicking in the field and entering the VLAN ID manually.

6 In the **Destinations** section, edit the destinations for the port mirroring session.

Depending on the type of port mirroring session being edited, different options are available for configuration.

Option	Description
Select a destination distributed port	Click the **Select distributed ports…** button to select ports from a list, or click the **Add distributed ports…** button to add ports by port number. You can add more than one distributed port.
Select a uplinks	Select an available uplink from the list and click **Add >** to add the uplink to the port mirroring session. You can select more than one uplink.
Select ports or uplinks	Click the **Select distributed ports…** button to select ports from a list, or click the **Add distributed ports…** button to add ports by port number. You can add more than one distributed port. Click the**Add uplinks...** button to add uplinks as the destination. Select uplinks from the list and click **OK**.
Specify IP address	Click the **Add** button. A new list entry is created. Select the entry and either click the Edit button to enter the IP address, or click directly into the IP Address field and enter the IP address. A warning dialog opens if the IP address is invalid.

7 Click **OK**.

Configure NetFlow Settings with the vSphere Web Client

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

NetFlow is available on a vSphere Distributed Switch version 5.0.0 and later. Version 5.1 and later of the switch supports IPFIX (NetFlow version 10).

Procedure

1 In the vSphere Web Client, navigate to the distributed switch.

2 Right-click the distributed switch in the navigator and select **All vCenter Actions > Edit Netflow**.

3 Type the **IP address** and **Port** of the NetFlow collector.

4 Type the **Switch 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.

5 (Optional) Set the **Active flow export timeout** and **Idle flow export timeout** in seconds.

6 (Optional) Set the **Sampling Rate**.

The sampling rate represents the number of packets that NetFlow drops after every collected packet. A sampling rate of *x* instructs NetFlow to drop packets in a *collected packets*:*dropped packets* ratio 1:*x*. If the rate is 0, NetFlow samples every packet, that is, collect one packet and drop none. If the rate is 1, NetFlow samples a packet and drops the next one, and so on.

7 (Optional) Enable or disable **Process internal flows only** with the drop-down menu.

When enabled, only data on network activity between virtual machines on the same host is collected.

8 Click **OK**.

Switch Discovery Protocol

Switch discovery protocols help vSphere administrators to determine which port of the physical switch is connected to a vSphere standard switch or vSphere distributed switch.

vSphere 5.0 and later 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 Web Client.

Enable Cisco Discovery Protocol on a vSphere Distributed Switch with the vSphere Web Client

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

Procedure

1 Browse to a distributed switch in the vSphere Web Client navigator.

2 Click the **Manage** tab, and click **Settings > Properties**.

3 Click **Edit**.

4 Click **Advanced**.

5 In the **Discovery Protocol** section, select **Cisco Discovery Protocol** from the **Type** drop-down menu .

6 Set the **Operation** from the drop-down menu.

Option	Description
Listen	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.
Advertise	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.
Both	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 in the vSphere Web Client

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 Web Client.

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

Procedure

1 Browse to a distributed switch in the vSphere Web Client navigator.

2 Click the **Manage** tab, and select **Settings > Properties**.

3 Click **Edit**.

4 Click **Advanced**.

5 Select **Link Layer Discovery Protocol** from the **Type** drop-down menu.

6 Set **Operation** to Listen, Advertise, or Both.

Operation	Description
Listen	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.
Advertise	ESXi makes information about the vSphere distributed switch available to the switch administrator, but does not detect and display information about the physical switch.
Both	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 with the vSphere Web Client

When CDP or LLDP is set to **Listen** or **Both**, you can view physical switch information from the vSphere Web Client.

Procedure

1 Browse to a host in the vSphere Web Client navigator.

2 Click the **Manage** tab, and click **Networking > Physical adapters**.

3 Select a physical adapter from the list to view detailed information.

Depending on the enabled switch discovery protocol, the properties of the switch appear under Cisco Discovery Protocol or Link Layer Discovery Protocol. If the information is available in the network, under Peer device capability you can see the system capabilities of the switch.

Mounting NFS Volumes

ESXi supports VMkernel-based NFS mounts for storing virtual disks on NFS datastores.

In addition to storing virtual disks on NFS datastores, you can also use NFS Datastores as a central repository for ISO images and virtual machine templates. For more information about creating NFS datastores, see the *vSphere Storage* documentation.

ESXi supports NFS version 3 over Layer 2 and Layer 3 Network switches.

Networking Rollback and Recovery

In vSphere 5.1 and later, you can prevent and recover from network misconfiguration by using rollback, backup configuration files or previous configurations.

vSphere 5.1 and later can rollback to previous valid networking configuration if a misconfiguration of the management network occurs. To recover from misconfiguration, you can connect directly to a host to fix any networking issues through the DCUI. Rollback is available for use on both standard and distributed switches.

vSphere Networking Rollback

By rolling configuration changes back, vSphere protects hosts from losing connection to vCenter Server as a result from misconfiguration of the management network.

In vSphere 5.1 and later, networking rollback is enabled by default. However, you can enable or disable rollbacks at the vCenter Server level.

Host Networking Rollbacks

Host networking rollbacks occur when an invalid change is made to the networking configuration for the connection with vCenter Server. Every network change that disconnects a host also triggers a rollback. The following examples of changes to the host networking configuration might trigger a rollback:

- Updating the speed or duplex of a physical NIC.
- Updating DNS and routing settings.
- Updating teaming and failover policies or traffic shaping policies of a standard port group that contains the management VMkernel network adapter.
- Updating the VLAN of a standard port group that contains the management VMkernel network adapter.

- Increasing the MTU of management VMkernel network adapter and its switch to values not supported by the physical infrastructure.
- Changing the IP settings of management VMkernel network adapters.
- Removing the management VMkernel network adapter from a standard or distributed switch.
- Removing a physical NIC of a standard or distributed switch containing the management VMkernel network adapter.

If a network disconnects for any of these reasons, the task fails and the host reverts to the last valid configuration.

vSphere Distributed Switch Rollbacks

Distributed switch rollbacks occur when invalid updates are made to distributed switches, distributed port groups, or distributed ports. The following changes to the distributed switch configuration trigger a rollback:

- Changing the MTU of a distributed switch.
- Changing the following settings in the distributed port group of the management VMkernel network adapter:
  - Teaming and failover
  - VLAN
  - Traffic shaping
- Blocking all ports in the distributed port group containing the management VMkernel network adapter.
- Overriding the policies on at the level of the distributed port for the management VMkernel network adapter.

If a configuration becomes invalid because of any of the changes, one or more hosts might become out of synchronization with the distributed switch.

If you know where the conflicting configuration setting is located, you can manually correct the setting. For example, if you have migrated a management VMkernel network adapter to a new VLAN, the VLAN might not be actually trunked on the physical switch. When you correct the physical switch configuration, the next distributed switch-to-host synchronization will resolve the configuration problem.

If you are not sure where the problem exists, you can restore the state of the distributed switch or distributed port group to an earlier configuration. See "Restore a vSphere Distributed Port Group Configuration with the vSphere Web Client," on page 46.

Disable Network Rollback with the vSphere Web Client

Rollback is enabled by default in vSphere 5.1 and later. You can disable rollback in vCenter Server using the vSphere Web Client.

Procedure

1 Browse to a vCenter Server instance in the vSphere Web Client navigator.

2 Click the **Manage** tab, and select **Settings**..

3 Select **Advanced Settings** and click **Edit**.

4 Select the `config.vpxd.network.rollback` key, and change the value to **false**.

If the key is not present, you can add it and set the value to false.

5 Click **OK**.

6 Restart vCenter Server to apply the changes.

Disable Network Rollback by Using the vCenter Server Configuration File

Rollback is enabled by default in vSphere 5.1 and later. You can disable rollback by editing the `vpxd.cfg` configuration file of vCenter Server directly.

Procedure

1 On the host machine of vCenter Server, navigate to the directory that contains the configuration file:

- On a Windows Server operating system, the location of the directory is *`vCenter Server home directory`*`\Application Data\VMware\VMware VirtualCenter`.
- On the vCenter Server Appliance, the location of the directory is `/etc/vmware-vpx`.

2 Open the `vpxd.cfg` file for editing.

3 In the `<network>` element, set the `<rollback>` element to **`false`**:

```
<config>
 <vpxd>
  <network>
    <rollback>false</rollback>
  </network>
 </vpxd>
</config>
```

4 Save and close the file.

5 Restart the vCenter Server system.

Restore a Previous Networking Configuration with the vSphere Web Client

You can restore a previous configuration of a vSphere distributed switch or port group to revoke invalid changes.

Procedure

1 In the vSphere Web Client, navigate to the affected a vSphere distributed switch, or to a distributed or uplink port group.

2 Right-click the affected switch or port group and select **All vCenter Actions > Restore Configuration**.

3 If you are restoring the configuration of a distributed switch, provide a backup file.

a Click **Browse** and navigate to the location of the distributed switch backup file.

b Select **Restore distributed switch and all port groups** or **Restore distributed switch only**.

c Click **Next**.

4 If you are restoring the configuration of a distributed port group or uplink port group:

a Select **Restore to previous configuration** or **Restore configuration from a file**.

Restore to previous configuration reverts the configuration one step back.

b If restoring from a file, click **Browse** to navigate to the location of a port group backup file.

c Click **Next**.

5 Review the configuration information and click **Finish**.

Resolve Errors in the Management Network Configuration on a vSphere Distributed Switch

In vSphere 5.1 and later, you can use the Direct Console User Interface (DCUI) to restore the connection between vCenter Server and a host that accesses the management network through a distributed switch.

If networking rollback is disabled, misconfiguring the port group for the management network on the distributed switch leads to loss of connection between vCenter Server and the hosts added to the switch. You have to use the DCUI to connect each host individually.

For more information about accessing and using the DCUI, see the *vSphere Security* documentation.

NOTE Recovery of the management connection on a distributed switch is not supported on stateless ESXi instances.

Prerequisites

Verify that the management network is configured on a port group on the distributed switch.

Procedure

1 Connect to the DCUI of the host.

2 From the **Network Restore Options** menu, select **Restore vDS**.

3 Configure the uplinks and optionally the VLAN for the management network.

4 Apply the configuration.

The DCUI creates a local ephemeral port and applies the values you provided for the VLAN and uplinks. The DCUI moves the VMkernel adapter for the management network to the new local port to restore connectivity to vCenter Server.

What to do next

After the connection of the host to vCenter Server is restored, correct the configuration of the distributed port group and re-add the VMkernel adapter to the group.

Configuring Protocol Profiles for Virtual Machine Networking

A network protocol profile contains a pool of IPv4 and IPv6 addresses that vCenter Server assigns to vApps or to virtual machines with vApp functionality that are connected to port groups associated with the profile.

Network protocol profiles also contain settings for the IP subnet, DNS, and HTTP proxy server.

To configure the networking settings of virtual machines by using from network protocol profiles, perform the following operations:

- Create network profiles at the level of a datacenter or a vSphere distributed switch.
- Associate a protocol profile with the port group of a vApp virtual machine.
- Enable the transient or static IP allocation policy from the settings of the vApp or from the vApp options of a virtual machine.

NOTE If you move a vApp or a virtual machine that retrieves its network settings from a protocol profile to another datacenter, to power it on you must assign a protocol profile to the connected port group on the destination datacenter.

- Add a Network Protocol Profile on page 165

  A network protocol profile contains a pool of IPv4 and IPv6 addresses. vCenter Server assigns those resources to vApps or to virtual machines with vApp functionality that are connected to port groups associated with the profile.

- Associate a Port Group with a Network Protocol Profile in the vSphere Web Client on page 167

  To apply the range of IP addresses from a network protocol profile to a virtual machine that is a part of a vApp or has vApp functionality enabled, associate the profile with a port group that controls the networking of the virtual machine.

- Configure a Virtual Machine or vApp to Use a Network Protocol Profile in the vSphere Web Client on page 168

  After you associate a protocol profile to a port group of a standard switch or a distributed switch, enable the usage of profile on a virtual machine that is connected to the port group and is associated with a vApp or has the vApp options enabled.

Add a Network Protocol Profile

A network protocol profile contains a pool of IPv4 and IPv6 addresses. vCenter Server assigns those resources to vApps or to virtual machines with vApp functionality that are connected to port groups associated with the profile.

Network protocol profiles also contain settings for the IP subnet, DNS, and HTTP proxy server.

NOTE If you move a vApp or a virtual machine that retrieves its network settings from a protocol profile to another datacenter, to power on the vApp or virtual machine you must assign a protocol profile to the connected port group on the destination datacenter.

Procedure

1 Navigate to a datacenter that is associated with the vApp and click the **Manage** tab.

2 Click **Network Protocol Profiles**

  Existing network protocol profiles are listed.

3 Click the Add icon (+) to add a new network protocol profile.

Select the Network Protocol Profile Name and Network

Name the network protocol profile and select the network that should use it.

Procedure

1 Type the name of the network protocol profile.

2 Select the networks that use this network protocol profile.

  A network can be associated with one network protocol profile at a time.

3 Click **Next**.

Specify Network Protocol Profile IPv4 Configuration

A network protocol profile contains a pool of IPv4 and IPv6 addresses for use by vApps. When you create a network protocol profile, you set up its IPv4 configuration.

You can configure network protocol profile ranges for IPv4, IPv6, or both. vCenter Server uses these ranges to dynamically allocate IP addresses to virtual machines when a vApp is set up to use transient IP allocation.

Procedure

1 Enter the **IP Subnet** and **Gateway** in their respective fields.

2 Select **DHCP Present** to indicate that the DHCP server is available on this network.

3 Enter the DNS server information.

Specify the servers by IP addresses separated by a comma, semicolon, or space.

4 Select the **Enable IP Pool** check box to specify an IP pool range.

5 If you enable IP Pools, enter a comma-separated list of host address ranges in the **IP pool range** field.

A range consists of an IP address, a pound sign (#), and a number indicating the length of the range.

The gateway and the ranges must be within the subnet. The ranges that you enter in the **IP pool range** field cannot include the gateway address.

For example, **`10.20.60.4#10, 10.20.61.0#2`** indicates that the IPv4 addresses can range from 10.20.60.4 to 10.20.60.13 and 10.20.61.0 to 10.20.61.1.

6 Click **Next**.

Specify Network Protocol Profile IPv6 Configuration

A network protocol profile contains a pool of IPv4 and IPv6 addresses for use by vApps. When you create a network protocol profile, you set up its IPv6 configuration.

You can configure network protocol profile ranges for IPv4, IPv6, or both.vCenter Server uses these ranges to dynamically allocate IP addresses to virtual machines when a vApp is set up to use transient IP allocation.

Procedure

1 Enter the **IP Subnet** and **Gateway** in their respective fields.

2 Select **DHCP Present** to indicate that the DHCP server is available on this network.

3 Enter the DNS server information.

Specify the servers by IP addresses separated by a comma, semicolon, or space.

4 Select the **Enable IP Pool** check box to specify an IP pool range.

5 If you enable IP Pools, enter a comma-separated list of host address ranges in the **IP pool range** field.

A range consists of an IP address, a pound sign (#), and a number indicating the length of the range. For example. assume that you specify the following IP pool range:

fe80:0:0:0:2bff:fe59:5a:2b#10,fe80:0:0:0:2bff:fe59:5f:b1#2

Then the addresses are in this range:

fe80:0:0:0:2bff:fe59:5a:2b - fe80:0:0:0:2bff:fe59:5a:34

and

fe80:0:0:0:2bff:fe59:5f:b1 - fe80:0:0:0:2bff:fe59:5f:b2

The gateway and the ranges must be within the subnet. The ranges that you enter in the **IP pool range** field cannot include the gateway address.

6 Click **Next**.

Specify Network Protocol Profile DNS and Other Configuration

When you create a network protocol profile, you can specify the DNS domain, DNS search path, a host prefix, and HTTP proxy.

Procedure

1 Enter the DNS domain.

2 Enter the host prefix.

3 Enter the DNS search path.

The search paths are specified as a list of DNS domains separated by commas, semi-colons, or spaces.

4 Enter the server name and port number for the proxy server.

The server name can optionally include a colon and a port number.

For example, `web-proxy:3912` is a valid proxy server.

5 Click **Next**.

Complete the Network Protocol Profile Creation

Procedure

- Review the settings and click **Finish** to complete adding the network protocol profile.

Associate a Port Group with a Network Protocol Profile in the vSphere Web Client

To apply the range of IP addresses from a network protocol profile to a virtual machine that is a part of a vApp or has vApp functionality enabled, associate the profile with a port group that controls the networking of the virtual machine.

You can associate a port group of a standard switch or a distributed port group of a distributed switch with a network protocol profile by using the settings of the group.

Procedure

1 Navigate to a distributed port group of a vSphere distributed switch or to a port group of a vSphere standard switch in the Networking view of the vSphere Web Client.

The port groups of standard switches are under the datacenter. The vSphere Web Client displays distributed port groups under the parent distributed switch object.

2 On the **Manage** tab, click **Network Protocol Profiles**.

3 Click **Associate a network protocol profile with the selected network**.

4 On the Set association type page of the Associate Network Protocol Profile wizard, select **Use an existing network protocol profile** and click **Next**.

If the existing network protocol profiles do not contain settings suitable for the vApp virtual machines in the port group, you must create a new profile.

5 Select the network protocol profile and click **Next**.

6 Examine the association and settings of the network protocol profile and click **Finish**.

Configure a Virtual Machine or vApp to Use a Network Protocol Profile in the vSphere Web Client

After you associate a protocol profile to a port group of a standard switch or a distributed switch, enable the usage of profile on a virtual machine that is connected to the port group and is associated with a vApp or has the vApp options enabled.

Prerequisites

Verify that the virtual machine is connected to a port group that is associated with the network protocol profile.

Procedure

1 Navigate to the virtual machine or the vApp in the inventory of the vSphere Web Client.

2 In the vSphere Web Client, open the settings of the vApp or the **vApp Options** tab of the virtual machine.

- Right-click a vApp and select **Edit settings**.
- Right-click a virtual machine, select **Edit settings**, and in the Edit Settings dialog box, click the **vApp Options** tab.

3 Click **Enable vApp options**.

4 Under Authoring, expand **IP allocation** and set the IP allocation scheme to **OVF environment**.

5 Under Deployment, expand **IP allocation** and set **IP allocation** to **Transient - IP Pool** or **Static - IP Pool**.

Both the **Static - IP Pool** and **Transient - IP Pool** options allocate an IP address from the range in the network protocol profile that is associated with the port group. If you select **Static - IP Pool**, the IP address is assigned the first time the virtual machine or vApp is powered on and the address persists across restarts. If you select **Transient - IP Pool**, an IP address is assigned every time the virtual machine or vApp is powered on.

6 Click **OK**.

When the virtual machine is powered on, the adapters connected to the port group receive IP addresses from the range in the protocol profile. When the virtual machine is powered off, the IP addresses are released.

Stateless Network Deployment

Stateless is a mode of execution for ESXi hosts with no local storage that formerly would save configuration or state. Configurations are abstracted into a host profile, which is a template that applies to a class of machines. Stateless allows easy replacement, removal, and addition of failed hardware, and improves the ease of scaling a hardware deployment.

Every stateless ESXi boot is like a first boot. The ESXi host boots with networking connectivity to vCenter Server through the built-in standard switch. If the host profile specifies distributed switch membership, vCenter Server joins the ESXi host to VMware distributed switches or a third party switch solution.

When planning the network setup for stateless ESXi hosts, you should keep the configuration as generic as possible and avoid host-specific items. Currently the design has no hooks to reconfigure physical switches when deploying a new host. Any such requirement would need special handling.

To set up stateless deployment, one ESXi host must be installed in the standard fashion. Then find and record the following network-related information to save in the host profile:

- vSphere standard switch instances and settings (port groups, uplinks, MTU, and so forth)
- Distributed switch instances (VMware and third party)
- Selection rules for uplinks and uplink port or port groups
- vNIC information:
  - Address information (IPv4 or IPv6, static or DHCP, gateway)
  - Port groups and distributed port groups assigned to the physical network adapter (`vmknic`)
  - If there are distributed switches, record VLAN, physical NICs bound to the `vmknic`, and if `Etherchannel` is configured

The recorded information is used as a template for the host profile. Once the host profile virtual switch information has been extracted and placed in the host profile, you have the opportunity to change any of the information. Modifications are offered for both standard and distributed switches in these sections: uplink selection policy, based on either vmnic name or device number, and auto discovery based on VLAN ID. The (possibly modified) information is stored by the stateless boot infrastructure and applied to a stateless ESXi host on its next boot. During network initialization, a generic network plug-in interprets the recorded host profile setting and does the following:

- Loads appropriate physical NIC drivers.
- Creates all standard switch instances, along with port groups. It selects uplinks based on policy. If the policy is based on the VLAN ID, there is a probing process to gather relevant information.
- For VMkernel network adapters connected to the standard switch, it creates VMkernel network adapters and connects them to port groups.
- For each VMkernel network adapter connected to a distributed switch, it creates a temporary standard switch (as needed) with uplinks bound to the VMkernel network adapter. It creates a temporary port group with VLAN and teaming policies based on recorded information. Specifically, IP-hash is used if Etherchannel was used in the distributed switch.
- Configures all VMkernel network adapter settings (assigns address, gateway, MTU, and so forth).

Basic connectivity is functioning, and the networking setup is complete if there is no distributed switch present.

If there is a distributed switch present, the system stays in maintenance mode until distributed switch remediation is complete. No virtual machines are started at this time. Because distributed switches requires vCenter Server, the boot process continues until vCenter Server connectivity is established, and vCenter Server notices that the host should be part of a distributed switch. It issues a distributed switch host join, creating a distributed switch proxy standard switch on the host, selects appropriate uplinks, and migrates the vmknic from the standard switch to the distributed switch. When this operation is complete, it deletes the temporary standard switch and port groups.

At the end of the remediation process, the ESXi host is taken out of maintenance mode, and HA or DRS can start virtual machines on the host.

In the absence of a host profile, a temporary standard switch is created with “default networking” logic, which creates a management network switch (with no VLAN tag) whose uplink corresponds to the PXE booting vNIC. A vmknic is created on the management network port group with the same MAC address as the PXE booting vNIC. This logic was previously used for PXE booting. If there is a host profile, but the networking host profile is disabled or fatally incomplete, vCenter Server falls back to default networking so that the ESXi host can be managed remotely. This triggers a compliance failure, so vCenter Server then initiates recovery actions.

Monitoring Network Packets 10

Monitor network packets that pass through the ports of a vSphere Standard Switch or a vSphere Distributed Switch to analyze the traffic between virtual machines and hosts.

Capturing and Tracing Network Packets by Using the pktcap-uw Utility

Monitor the traffic that flows through physical network adapters, VMkernel adapters, and virtual machines adapters, and analyze packet information by using the graphical user interface of network analysis tools such as Wireshark.

In vSphere 5.5 or later, you can monitor packets on a host by using the `pktcap-uw` console utility. You can use the utility without additional installation on an ESXi host. `pktcap-uw` provides many points in the host network stack at which you can monitor traffic.

For detailed analysis of captured packets, you can save packet content from the `pktcap-uw` utility to files in PCAP or PCAPNG format and open them in Wireshark. You can also troubleshoot dropped packets and trace a packet's path in the network stack.

Note The `pktcap-uw` utility is not fully supported for backward compatibility across vSphere releases. The options of the utility might change in the future.

pktcap-uw Command Syntax for Capturing Packets

Use the `pktcap-uw` utility to inspect the contents of packets while they traverse the network stack on an ESXi host.

pktcap-uw Syntax for Capturing Packets

The `pktcap-uw` command has the following syntax for capturing packets at a certain place in the network stack:

```
pktcap-uw switch_port_arguments capture_point_options filter_options output_control_options
```

Note Certain options of the `pktcap-uw` utility are designed for VMware internal use only and you should use them only under the supervision of VMware Technical Support. These options are not described in the *vSphere Networking* guide.

Table 10-1. pktcap-uw Arguments for Capturing Packets

Argument Group	Argument	Description
switch_port_arguments	`--uplink vmnic`*X*	Capture packets that are related to a physical adapter. You can combine the `--uplink` and `--capture` options for monitoring packets at a certain place in the path between the physical adapter and the virtual switch. See "Capture Packets That Arrive at a Physical Adapter," on page 175.
	`--vmk vmk`*X*	Capture packets that are related to a VMKernel adapter. You can combine the `vmk` and `--capture` options for monitoring packets at a certain place in the path between the VMkernel adapter and the virtual switch. See "Capture Packets for a VMkernel Adapter," on page 178.
	`--switchport {`*vmxnet3_port_ID* `|` *vmkernel_adapter_port_ID*`}`	Capture packets that are related to a VMXNET3 virtual machine adapter or to a VMkernel adapter that is connected to a particular virtual switch port. You can view the ID of the port in the network panel of the `esxtop` utility. You can combine the `switchport` and `capture` options for monitoring packets at a certain place in the path between the VMXNET3 adapter or VMkernel adapter and the virtual switch. See "Capture Packets for a VMXNET3 Virtual Machine Adapter," on page 177.
	`--lifID` *lif_ID*	Capture packets that are related to the logical interface of a distributed router. See the *VMware NSX* documentation.
capture_point_options	`--capture` *capture_point*	Capture packets at a particular place in the network stack. For example, you can monitor packets right after they arrive from a physical adapter.

Table 10-1. pktcap-uw Arguments for Capturing Packets (Continued)

Argument Group	Argument	Description
	`--dir {0|1}`	Capture packets according to the direction of the flow with regard to the virtual switch. 0 stands for incoming traffic and 1 for outgoing traffic. By default, the `pktcap-uw` utility captures ingress traffic. Use the `--dir` option together with the `--uplink`, `--vmk`, or `--switchport` option.
	`--stage {0|1}`	Capture the packet closer to its source or to its destination. Use this option to examine how a package changes while it traverses the points in the stack. 0 stands for traffic closer to source and 1 for traffic closer to destination. Use the `--stage` option together with the `--uplink`, `--vmk` , `--switchport`, or `--dvfilter` option.
	`--dvfilter` *`filter_name`* `--capture PreDVFilter|PostDVFilter`	Capture packets before or after a vSphere Network Appliance (DVFilter) intercepts them. See "Capture Packets at DVFilter Level," on page 180.
	`-A | --availpoints`	View all capture points that the `pktcap-uw` utility supports.
	For details about the capture points of the `pktcap-uw` utility, see "Capture Points of the pktcap-uw Utility," on page 182.	
filter_options	Filter captured packets according to source or destination address, VLAN ID, VXLAN ID, Layer 3 protocol, and TCP port. See "pktcap-uw Options for Filtering Packets," on page 174.	
output_control_options	Save the contents of a packet to a file, capture only a number of packets, and capture a number of bytes at the beginning of packets, and so on. See "pktcap-uw Options for Output Control," on page 174.	

The vertical bars | represent alternative values, and the curly brackets {} used with vertical bars specify a list of choices for an argument or option.

pktcap-uw Command Syntax for Tracing Packets

Use the `pktcap-uw` utility to view the path of a packet in the network stack on an ESXi host for latency analysis.

pktcap-uw Syntax for Tracing Packets

The command of the `pktcap-uw` utility has the following syntax for tracing packets in the network stack:

`pktcap-uw --trace` *`filter_options output_control_options`*

Options to the pktcap-uw Utility for Tracing Packets

The `pktcap-uw` utility supports the following options when you use it to trace packets:

Table 10-2. pktcap-uw Options for Tracing Packets

Argument	Description
filter_options	Filter traced packets according to source or destination address, VLAN ID, VXLAN ID, Layer 3 protocol, and TCP port. See "pktcap-uw Options for Filtering Packets," on page 174.
output_control_options	Save the contents of a packet to a file and trace only a number of packets. See "pktcap-uw Options for Output Control," on page 174.

pktcap-uw Options for Output Control

Use the options for output control of the `pktcap-uw` utility to save packet contents to a file, capture up to a certain number of bytes from each packet, and limit the number of captured packets.

pktcap-uw Options for Output Control

The options of the `pktcap-uw` utility for output control are valid when you capture and trace packets. For information about the command syntax of the `pktcap-uw` utility, see "pktcap-uw Command Syntax for Capturing Packets," on page 171 and "pktcap-uw Command Syntax for Tracing Packets," on page 173.

Table 10-3. Options for Output Control That Are Supported by the pktcap-uw Utility

Option	Description
`{-o | --outfile}` *`pcap_file`*	Save captured or traced packets to a file in packet capture (PCAP) format. Use this option to examine packets in a visual analyzer tool such as Wireshark.
`-P | --ng`	Save packet content in the PCAPNG file format. Use this option together with the `-o` or `--outfile` option.
`--console`	Print packet details and content to the console output. By default, the `pktcap-uw` utility shows packet information in the console output.
`{-c | --count}` *`number_of_packets`*	Capture the first *number_of_packets* packets.
`{-s | --snaplen}` *`snapshot_length`*	Capture only the first *snapshot_length* bytes from each packet. If traffic on the host is intensive, use this option to reduce the load on the CPU and storage. To limit the size of captured contents, set a value greater than 24. To capture the complete packet, set this option to 0.
`-h`	View help about the `pktcap-uw` utility.

The vertical bars | represent alternative values, and the curly brackets {} used with vertical bars specify a list of choices for an argument or option.

pktcap-uw Options for Filtering Packets

Narrow the range of packets that you monitor by using the `pktcap-uw` utility to apply filtering options for source and destination address, VLAN, VXLAN, and next level protocol consuming the packet payload.

Filter Options

The filter options for `pktcap-uw` are valid when you capture and trace packets. For information about the command syntax of the `pktcap-uw` utility, see "pktcap-uw Command Syntax for Capturing Packets," on page 171 and "pktcap-uw Command Syntax for Tracing Packets," on page 173.

Table 10-4. Filter Options of the pktcap-uw Utility

Option	Description
`--srcmac` *mac_address*	Capture or trace packets that have a specific source MAC address. Use colons to separate the octets in it.
`--dstmac` *mac_address*	Capture or trace packets that have a specific destination MAC address. Use colons to separate the octets in it.
`--mac` *mac_address*	Capture or trace packets that have a specific source or destination MAC address. Use colons to separate the octets in it.
`--ethtype 0x`*Ethertype*	Capture or trace packets at Layer 2 according to the next level protocol that consumes packet payload. *EtherType* corresponds to the EtherType field in Ethernet frames . It represents the type of next level protocol that consumes the payload of the frame. For example, to monitor traffic for the Link Layer Discovery Protocol (LLDP) protocol, type **`--ethtype 0x88CC`**.
`--vlan` *VLAN_ID*	Capture or trace packets that belong to a VLAN.
`--srcip` *IP_addess*|*IP_address/subnet_range*	Capture or trace packets that have a specific source IPv4 address or subnet.
`--dstip` *IP_addess*|*IP_address/subnet_range*	Capture or trace packets that have a specific destination IPv4 address or subnet.
`--ip` *IP_addess*	Capture or trace packets that have a specific source or destination IPv4 address.
`--proto 0x`*IP_protocol_number*	Capture or trace packets at Layer 3 according to the next level protocol that consumes the payload. For example, to monitor traffic for the UDP protocol, type **`--proto 0x11`**.
`--srcport` *source_port*	Capture or trace packets according to their source TCP port.
`--dstport` *destination_port*	Capture or trace packets according to their destination TCP port.
`--tcpport` *TCP_port*	Capture or trace packets according to their source or destination TCP port.
`--vxlan` *VXLAN_ID*	Capture or trace packets that belong to a VXLAN.

The vertical bars | represent alternative values.

Capturing Packets by Using the pktcap-uw Utility

Capture packets through the `pktcap-uw` utility in the path between a virtual switch and the physical adapters, VMkernel adapters and virtual machine adapters to troubleshoot data transfer in the network stack on an ESXi host.

Capture Packets That Arrive at a Physical Adapter

Monitor host traffic related to the external network by capturing packets at certain points in the path between a vSphere Standard Switch or vSphere Distributed Switch and a physical adapter.

You can specify a certain capture point in the data path between a virtual switch and a physical adapter, or determine a capture point by traffic direction with regard to the switch and proximity to the packet source or destination. For information about supported capture points, see "Capture Points of the pktcap-uw Utility," on page 182.

Procedure

1 (Optional) Find the name of the physical adapter that you want to monitor in the host adapter list.

- In the vSphere Web Client, on the **Manage** tab for the host, click **Networking** and select **Physical adapters**.
- In the ESXi Shell to the host, to view a list of the physical adapters and examine their state, run the following ESXCLI command:

  ```
  esxcli network nic list
  ```

Each physical adapter is represented as `vmnic`*X*. *X* is the number that ESXi assigned to the physical adapter port.

2 In the ESXi Shell to the host, run the `pktcap-uw` command with the `--uplink vmnic`*X* argument and with options to monitor packets at a particular point, filter captured packets and save the result to a file.

`pktcap-uw --uplink vmnic`*X* `[--capture` *capture_point*`|--dir 0|1] [`*filter_options*`] [--outfile` *pcap_file_path* `[--ng]] [--count` *number_of_packets*`]`

where the square brackets [] enclose the options of the `pktcap-uw --uplink vmnic`*X* command and the vertical bars | represent alternative values.

If you run the `pktcap-uw --uplink vmnic`*X* command without options, you obtain the content of packets that are incoming to the standard or distributed switch in the console output at the point where they are switched.

a Use the `--capture` option to check packets at another capture point or the `--dir` option at another traffic direction.

pktcap-uw Command Option	Goal
`--capture UplinkSnd`	Monitor packets immediately before they enter the physical adapter device.
`--capture UplinkRcv`	Monitor packets immediately after they are received in the network stack from the physical adapter.
`--dir 1`	Monitor packets that leave the virtual switch.
`--dir 0`	Monitor packets that enter the virtual switch.

b Use a *filter_options* to filter packets according to source and destination address, VLAN ID, VXLAN ID, Layer 3 protocol, and TCP port.

For example, to monitor packets from a source system that has IP address 192.168.25.113, use the `--srcip 192.168.25.113` filter option.

c Use options to save the contents of each packet or the contents of a limited number of packets to a `.pcap` or `.pcapng` file.

- To save packets to a `.pcap` file, use the `--outfile` option.
- To save packets to a `.pcapng` file, use the `--ng` and `--outfile` options.

You can open the file in a network analyzer tool such as Wireshark.

By default, the `pktcap-uw` utility saves the packet files to the root folder of the ESXi file system.

d Use the `--count` option to monitor only a number of packets.

3 If you have not limited the number of packets by using the `--count` option, press Ctrl+C to stop capturing or tracing packets.

Example: Capture Packets That Are Received at vmnic0 from an IP Address 192.168.25.113

To capture the first 60 packets from a source system that is assigned the IP address 192.168.25.113 at vmnic0 and save them to a file called `vmnic0_rcv_srcip.pcap`, run the following `pktcap-uw` command:

```
pktcap-uw --uplink vmnic0 --capture UplinkRcv --srcip 192.168.25.113 --outfile
vmnic0_rcv_srcip.pcap --count 60
```

What to do next

If the contents of the packet are saved to a file, copy the file from the ESXi host to the system that runs a graphical analyzer tool, such as Wireshark, and open it in the tool to examine the packet details.

Capture Packets for a VMXNET3 Virtual Machine Adapter

Monitor traffic that flows between a virtual switch and a VMXNET3 virtual machine adapter by using the `pktcap-uw` utility.

You can specify a certain capture point in the data path between a virtual switch and a virtual machine adapter. You can also determine a capture point by traffic direction with regard to the switch and proximity to the packet source or destination. For information about supported capture points, see "Capture Points of the pktcap-uw Utility," on page 182.

Prerequisites

Verify that the virtual machine adapter is of type VMXNET3.

Procedure

1 On the host, learn the port ID of the virtual machine adapter by using the `esxtop` utility.

   a In the ESXi Shell to the host, to start the utility, run `esxtop`.

   b Press N to switch to the network panel of the utility.

   c In the USED-BY column, locate the virtual machine adapter, and write down the PORT-ID value for it.

      The USED-BY field contains the name of the virtual machine and the port to which the virtual machine adapter is connected.

   d Press Q to exit `esxtop`.

2 In the ESXi Shell to the host, run `pktcap-uw --switchport` *`port_ID`*.

   port_ID is the ID that the `esxtop` utility displays for the virtual machine adapter in the PORT-ID column.

3 In the ESXi Shell to the host, run the `pktcap-uw` command with the `--switchport` *`port_ID`* argument and with options to monitor packets at a particular point, filter captured packets and save the result to a file.

   ```
   pktcap-uw --switchport port_ID [--capture capture_point|--dir 0|1 --stage 0|1]
   [filter_options] [--outfile pcap_file_path [--ng]] [--count number_of_packets]
   ```

   where the square brackets [] enclose the options of the `pktcap-uw --switchport` *`port_ID`* command and the vertical bars | represent alternative values.

If you run the `pktcap-uw --switchport port_ID` command without options, you obtain the content of packets that are incoming to the standard or distributed switch in the console output at the point when they are switched.

a To check packets at another capture point or direction in the path between the guest operating system and the virtual switch, use the `--capture` option or combine the values of the `--dir` and `--stage` options.

pktcap-uw Command Options	Goal
`--capture Vmxnet3Tx`	Monitor packets when they pass from the virtual machine to the switch.
`--capture Vmxnet3Rx`	Monitor packets when they arrive to the virtual machine.
`--dir 1 --stage 0`	Monitor packets immediately after they leave the virtual switch.
`--dir 1`	Monitor packets immediately before they enter the virtual machine.
`--dir 0 --stage 1`	Monitor packets immediately after they enter the virtual switch.

b Use a *filter_options* to filter packets according to source and destination address, VLAN ID, VXLAN ID, Layer 3 protocol, and TCP port.

For example, to monitor packets from a source system that has IP address 192.168.25.113, use the `--srcip 192.168.25.113` filter option.

c Use options to save the contents of each packet or the contents of a limited number of packets to a `.pcap` or `.pcapng` file.

- To save packets to a `.pcap` file, use the `--outfile` option.
- To save packets to a `.pcapng` file, use the `--ng` and `--outfile` options.

You can open the file in a network analyzer tool such as Wireshark.

By default, the `pktcap-uw` utility saves the packet files to the root folder of the ESXi file system.

d Use the `--count` option to monitor only a number of packets.

4 If you have not limited the number of packets by using the `--count` option, press Ctrl+C to stop capturing or tracing packets.

Example: Capture Packets That Are Received at a Virtual Machine from an IP Address 192.168.25.113

To capture the first 60 packets from a source that is assigned the IP address 192.168.25.113 when they arrive at a virtual machine adapter with port ID 33554481 and save them to a file called `vmxnet3_rcv_srcip.pcap`, run the following `pktcap-uw` command:

```
pktcap-uw --switchport 33554481 --capture Vmxnet3Rx --srcip 192.168.25.113 --outfile
vmxnet3_rcv_srcip.pcap --count 60
```

What to do next

If the contents of the packet are saved to a file, copy the file from the ESXi host to the system that runs a graphical analyzer tool, such as Wireshark, and open it in the tool to examine the packet details.

Capture Packets for a VMkernel Adapter

Monitor packets that are exchanged between a VMkernel adapter and a virtual switch by using the `pktcap-uw` utility.

You can capture packets at a certain capture point in the flow between a virtual switch and a VMkernel adapter. You can also determine a capture point by traffic direction with regard to the switch and proximity to the packet source or destination. For information about supported capture points, see "Capture Points of the pktcap-uw Utility," on page 182.

Procedure

1 (Optional) Find the name of the VMkernel adapter that you want to monitor in the VMkernel adapter list.

- In the vSphere Web Client, from the Networking list on the **Manage** tab for the host, select **VMkernel adapters**.
- In the ESXi Shell to the host, to view a list of the physical adapters, run the following console command:

  ```
  esxcli network ip interface list
  ```

Each VMkernel adapter is represented as `vmkX`, where *X* is the sequence number that ESXi assigned to the adapter.

2 In the ESXi Shell to the host, run the `pktcap-uw` command with the `--vmk vmkX` argument and with options to monitor packets at a particular point, filter captured packets and save the result to a file.

```
pktcap-uw --vmk vmkX [--capture capture_point|--dir 0|1 --stage 0|1]  [filter_options]
[--outfile pcap_file_path [--ng]] [--count number_of_packets]
```

where the square brackets [] enclose the options of the `pktcap-uw --vmk vmkX` command and the vertical bars | represent alternative values.

You can replace the `--vmk vmkX` option with `--switchport vmkernel_adapter_port_ID`, where *vmkernel_adapter_port_ID* is the PORT-ID value that the network panel of the `esxtop` utility displays for the adapter.

If you run the `pktcap-uw --vmk vmkX` command without options, you obtain the content of packets that are leaving the VMkernel adapter.

a To check transmitted or received packets at a specific place and direction, use the `--capture` option, or combine the values of the `--dir` and `--stage` options.

pktcap-uw Command Options	Goal
`--dir 1 --stage 0`	Monitor packets immediately after they leave the virtual switch.
`--dir 1`	Monitor packets immediately before they enter the VMkernel adapter.
`--dir 0 --stage 1`	Monitor packets immediately before they enter the virtual switch.

b Use a *filter_options* to filter packets according to source and destination address, VLAN ID, VXLAN ID, Layer 3 protocol, and TCP port.

For example, to monitor packets from a source system that has IP address 192.168.25.113, use the `--srcip 192.168.25.113` filter option.

c Use options to save the contents of each packet or the contents of a limited number of packets to a `.pcap` or `.pcapng` file.

- To save packets to a `.pcap` file, use the `--outfile` option.
- To save packets to a `.pcapng` file, use the `--ng` and `--outfile` options.

You can open the file in a network analyzer tool such as Wireshark.

By default, the `pktcap-uw` utility saves the packet files to the root folder of the ESXi file system.

d Use the `--count` option to monitor only a number of packets.

3 If you have not limited the number of packets by using the `--count` option, press Ctrl+C to stop capturing or tracing packets.

What to do next

If the contents of the packet are saved to a file, copy the file from the ESXi host to the system that runs a graphical analyzer tool, such as Wireshark, and open it in the tool to examine the packet details.

Capture Dropped Packets

Troubleshoot lost connectivity by capturing dropped packets through the `pktcap-uw` utility.

A packet might be dropped at a point in the network stream for many reasons, for example, a firewall rule, filtering in an IOChain and DVfilter, VLAN mismatch, physical adapter malfunction, checksum failure, and so on. You can use the `pktcap-uw` utility to examine where packets are dropped and the reason for the drop.

Procedure

1 In the ESXi Shell to the host, run the `pktcap-uw --capture Drop` command with options to monitor packets at a particular point, filter captured packets and save the result to a file.

`pktcap-uw --capture Drop [`*`filter_options`*`] [--outfile `*`pcap_file_path`*` [--ng]] [--count `*`number_of_packets`*`]`

where the square brackets [] enclose the options of the `pktcap-uw --capture Drop` command and the vertical bars | represent alternative values.

a Use a *filter_options* to filter packets according to source and destination address, VLAN ID, VXLAN ID, Layer 3 protocol, and TCP port.

For example, to monitor packets from a source system that has IP address 192.168.25.113, use the `--srcip 192.168.25.113` filter option.

b Use options to save the contents of each packet or the contents of a limited number of packets to a `.pcap` or `.pcapng` file.

- To save packets to a `.pcap` file, use the `--outfile` option.
- To save packets to a `.pcapng` file, use the `--ng` and `--outfile` options.

You can open the file in a network analyzer tool such as Wireshark.

By default, the `pktcap-uw` utility saves the packet files to the root folder of the ESXi file system.

NOTE You can see the reason and the place where a packet is dropped only when you capture packets to the console output. The `pktcap-uw` utility saves only the content of packets to a `.pcap` or `.pcapng` file.

c Use the `--count` option to monitor only a number of packets.

2 If you have not limited the number of packets by using the `--count` option, press Ctrl+C to stop capturing or tracing packets.

Besides the contents of dropped packets, the output of the `pktcap-uw` utility displays the reason for the drop and the function in the network stack that handled the packet last.

What to do next

If the contents of the packet are saved to a file, copy the file from the ESXi host to the system that runs a graphical analyzer tool, such as Wireshark, and open it in the tool to examine the packet details.

Capture Packets at DVFilter Level

Examine how packets change when they pass through a vSphere Network Appliance (DVFilter).

DVFilters are agents that reside in the stream between a virtual machine adapter and a virtual switch. They intercept packets to protect virtual machines from security attacks and unwanted traffic.

Procedure

1 (Optional) To find the name of the DVFilter that you want to monitor, in the ESXi Shell, run the `summarize-dvfilter` command.

The output of the command contains the fast-path and slow-path agents of the DVFilters that are deployed on the host.

2 Run the `pktcap-uw` utility with the `--dvfilter` *`dvfilter_name`* argument and with options to monitor packets at a particular point, filter captured packets and save the result to a file.

`pktcap-uw --dvFilter` *`dvfilter_name`* `--capture PreDVFilter|PostDVFilter [`*`filter_options`*`] [--outfile` *`pcap_file_path`* `[--ng]] [--count` *`number_of_packets`*`]`

where the square brackets [] enclose optional items of the `pktcap-uw --dvFilter vmnic`*`X`* command and the vertical bars | represent alternative values.

a Use the `--capture` option to monitor packets before or after the DVFilter intercepts them.

pktcap-uw Command Option	Goal
`--capture PreDVFilter`	Capture packets before they enter the DVFilter.
`--capture PostDVFilter`	Capture packets after they leave the DVFilter.

b Use a *filter_options* to filter packets according to source and destination address, VLAN ID, VXLAN ID, Layer 3 protocol, and TCP port.

For example, to monitor packets from a source system that has IP address 192.168.25.113, use the `--srcip 192.168.25.113` filter option.

c Use options to save the contents of each packet or the contents of a limited number of packets to a `.pcap` or `.pcapng` file.

- To save packets to a `.pcap` file, use the `--outfile` option.
- To save packets to a `.pcapng` file, use the `--ng` and `--outfile` options.

You can open the file in a network analyzer tool such as Wireshark.

By default, the `pktcap-uw` utility saves the packet files to the root folder of the ESXi file system.

d Use the `--count` option to monitor only a number of packets.

3 If you have not limited the number of packets by using the `--count` option, press Ctrl+C to stop capturing or tracing packets.

What to do next

If the contents of the packet are saved to a file, copy the file from the ESXi host to the system that runs a graphical analyzer tool, such as Wireshark, and open it in the tool to examine the packet details.

Using the Capture Points of the pktcap-uw Utility

You use the capture points of the `pktcap-uw` utility to monitor packets when a function handles them at a specific place in the network stack on a host.

Overview of Capture Points

A capture point in the `pktcap-uw` utility represents a place in the path between a virtual switch on one side and a physical adapter, VMkernel adapter or a virtual machine adapter on the other.

You can use certain capture points in combination with an adapter option. For example, you use the UplinkRcv point when you capture uplink traffic. You can address other points standalone. For example, use the Drop point to inspect all dropped packets.

NOTE Certain capture points of the `pktcap-uw` utility are designed for VMware internal use only and you should use them only under the supervision of VMware Technical Support. These capture points are not described in the *vSphere Networking* guide.

Option for Using Capture Points in the pktcap-uw Utility

To examine a packet state or content at a capture point, add the `--capture`*capture_point* option to the `pktcap-uw` utility.

Auto-Selecting a Capture Point

For traffic that is related to a physical, VMkernel or VMXNET3 adapter, by combining the `--dir` and `--stage` options you can auto-select and switch between capture points to examine how a packet changes before and after a point.

Capture Points of the pktcap-uw Utility

The `pktcap-uw` utility supports capture points that can be used only when you monitor uplink, VMkernel or virtual machine traffic, and capture points that represent special places in the stack that are not related to the adapter type.

Capture Points That Are Relevant to Physical Adapter Traffic

The `pktcap-uw --uplink vmnic`*X* command supports capture points for functions that handle traffic at a specific place and direction in the path between the physical adapter and the virtual switch.

Capture Point	Description
UplinkRcv	The function that receives packets from the physical adapter.
UplinkSnd	The function that sends packets to the physical adapter.
PortInput	The function that passes a list of packets from UplinkRcv to a port on the virtual switch.
PortOutput	The function that passes a list of packets from a port on the virtual switch to the UplinkSnd point.

Capture Points That Are Relevant to Virtual Machine Traffic

The `pktcap-uw --switchport` *vmxnet3_port_ID* command supports capture points for functions that handle traffic packets at a specific place and direction in the path between a VMXNET3 adapter and a virtual switch.

Capture Point	Description
Vmxnet3Rx	The function in the VMXNET3 backend that receives packets from the virtual switch.
Vmxnet3Tx	The function in the VMXNET3 backend that sends packets from the virtual machine to the virtual switch.
PortOutput	The function that passes a list of packets from a port on the virtual switch to Vmxnet3Rx.
PortInput	The function that passes a list of packets from Vmxnet3Tx to a port on the virtual switch. Default capture point for traffic related to a VMXNET3 adapter.

Capture Points That Are Relevant to VMkernel Adapter Traffic

The `pktcap-uw --vmk vmk`*X* and `pktcap-uw --switchport` *vmkernel_adapter_port_ID* commands support capture points that represent functions at a specific place and direction in the path between a VMkernel adapter and a virtual switch.

Capture Point	Description
PortOutput	The function that passes a list of packets from a port on the virtual switch to the VMkernel adapter.
PortInput	The function that passes a list of packets from the VMkernel adapter to a port on the virtual switch. Default capture point for traffic related to a VMkernel adapter.

Capture Points That Are Relevant to Distributed Virtual Filters

The `pktcap-uw --dvfilter divfilter_name` command requires a capture point that indicates whether to capture packets when they enter the DVFilter or when they leave it.

Capture Point	Description
PreDVFilter	The point before a DVFilter intercepts a packet.
PostDVFilter	The point after a DVFilter intercepts a packet.

Standalone Capture Points

Certain capture points are mapped directly to the network stack rather than to a physical, VMkernel or VMXNET3 adapter.

Capture Point	Description
Drop	Captures dropped packets and shows the place where drops occur.
TcpipDispatch	Capture packets at the function that dispatches traffic to the TCP/IP stack of the VMkernel from the virtual switch, and the reverse.
PktFree	Capture packets right before they are released.
VdrRxLeaf	Capture packets at the receive leaf I/O chain of a dynamic router in VMware NSX. Use this capture point together with the `--lifID` option.
VdrRxTerminal	Capture packets at the receive terminal I/O chain of a dynamic router in VMware NSX. Use this capture point together with the `--lifID` option.
VdrTxLeaf	Capture packets at the transmit leaf I/O chain of a dynamic router in VMware NSX. Use this capture point together with the `--lifID` option.
VdrTxTerminal	Capture packets at the transmit terminal I/O chain of a dynamic router in VMware NSX. Use this capture point together with the `--lifID` option.

For information about dynamic routers, see the *VMware NSX* documentation.

List the Capture Points of the pktcap-uw Utility

View all capture points of the `pktcap-uw` utility to find the name of capture point for monitoring traffic at a certain place in the network stack on the ESXi host.

For information about the capture points of the `pktcap-uw` utility, see "Capture Points of the pktcap-uw Utility," on page 182.

Procedure

- In the ESXi Shell to the host, run the `pktcap-uw -A` command to view all capture points that the `pktcap-uw` utility supports.

Trace Packets by Using the pktcap-uw Utility

Use the `pktcap-uw` utility to trace the path that packets traverse in the network stack for latency analysis and for locating the point where a packet is corrupted or dropped.

The `pktcap-uw` utility shows the path of packets together with timestamps that note the time when a packet is handled by a networking function on ESXi. The utility reports the path of a packet immediately before it is released from the stack.

To view the full path information for a packet, you must print the result from the `pktcap-uw` utility in the console output or save it to a PCAPNG file.

Procedure

1 In the ESXi Shell to the host, run the `pktcap-uw --trace` command with options to filter traced packets, save the result to a file and limit the number of traced packets.

`pktcap-uw --trace [filter_options] [--outfile pcap_file_path [--ng]] [--count number_of_packets]`

where the square brackets [] enclose optional items of the `pktcap-uw --trace` command and the vertical bars | represent alternative values.

a Use a *filter_options* to filter packets according to source and destination address, VLAN ID, VXLAN ID, Layer 3 protocol, and TCP port.

For example, to monitor packets from a source system that has IP address 192.168.25.113, use the `--srcip 192.168.25.113` filter option.

b Use options to save the contents of each packet or the contents of a limited number of packets to a `.pcap` or `.pcapng` file.

- To save packets to a `.pcap` file, use the `--outfile` option.
- To save packets to a `.pcapng` file, use the `--ng` and `--outfile` options.

You can open the file in a network analyzer tool such as Wireshark.

By default, the `pktcap-uw` utility saves the packet files to the root folder of the ESXi file system.

NOTE A `.pcap` file contains only the contents of traced packets. To collect packet paths besides packet content, save the output to a `.pcapng` file.

c Use the `--count` option to monitor only a number of packets.

2 If you have not limited the number of packets by using the `--count` option, press Ctrl+C to stop capturing or tracing packets.

What to do next

If the contents of the packet are saved to a file, copy the file from the ESXi host to the system that runs a graphical analyzer tool, such as Wireshark, and open it in the tool to examine the packet details.

11 Networking Best Practices

Consider these best practices when you configure your network.

- Isolate from one another the networks for host management, vSphere vMotion, vSphere FT, and so on, to improve security and performance.

  Assign a group 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 handle application traffic, for example, from a Web client.

- 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, verify with your network administrator that the networks or VLANs you choose are isolated from the rest of your environment and that no routers connect them.
- Keep the vSphere vMotion connection on a separate network. 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.
- 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 VMXNET 3 virtual machine NICs.
- Physical network adapters 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 in a vSphere Distributed Switch with the same MTU. When several VMkernel network adapters, configured with different MTUs, are connected to vSphere distributed switches, you might experience network connectivity problems.
- When creating a distributed port group, do not use dynamic port binding. Dynamic port binding has been deprecated since ESXi 5.0.

Index

Q

R

S

T

U

V