



Virtual Switch Policies

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

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

The vSwitch policies consist of:

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

Layer 2 Security Policy

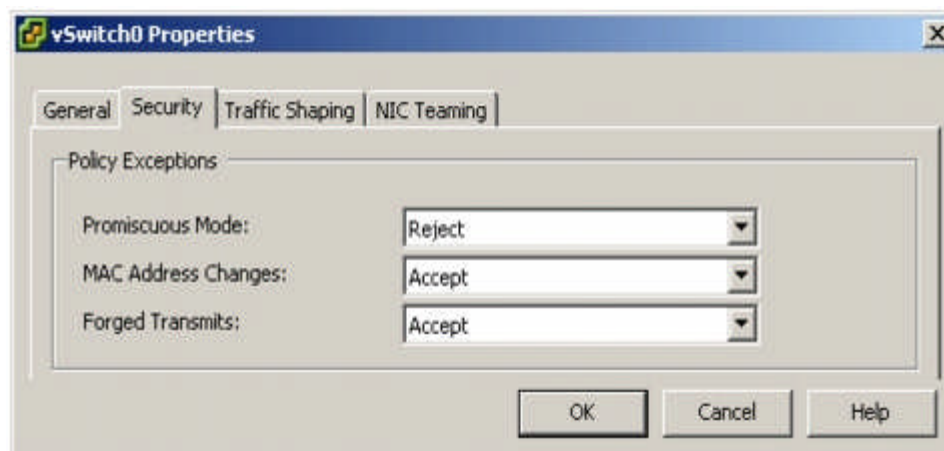
Layer 2 is the data link layer. The three elements of the Layer 2 Security policy are promiscuous mode, MAC address changes, and forged transmits.

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

For further information on security, see [Securing Virtual Switch Ports](#).

To edit the Layer 2 Security policy

- 1 Log into the VMware VI Client and select the server from the inventory panel.
The hardware configuration page for this server appears.
- 2 Click the **Configuration** tab, and click **Networking**.
- 3 Click **Properties** for the vSwitch whose Layer 2 Security policy you want to edit.
- 4 In the **Properties** dialog box for the vSwitch, click the **Ports** tab.
- 5 Select the vSwitch item and click **Edit**.
- 6 In the Properties dialog box for the vSwitch, click the **Security** tab.



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

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

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

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

If the Guest OS changes the MAC address back to match the MAC address in the .vmx configuration file, inbound frames will be 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 will be dropped.
 - **Accept** — No filtering is performed and all outbound frames are passed.

8 Click **OK**.

Traffic Shaping Policy

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

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

To edit the Traffic Shaping policy

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

The hardware configuration page for this server appears.

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

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

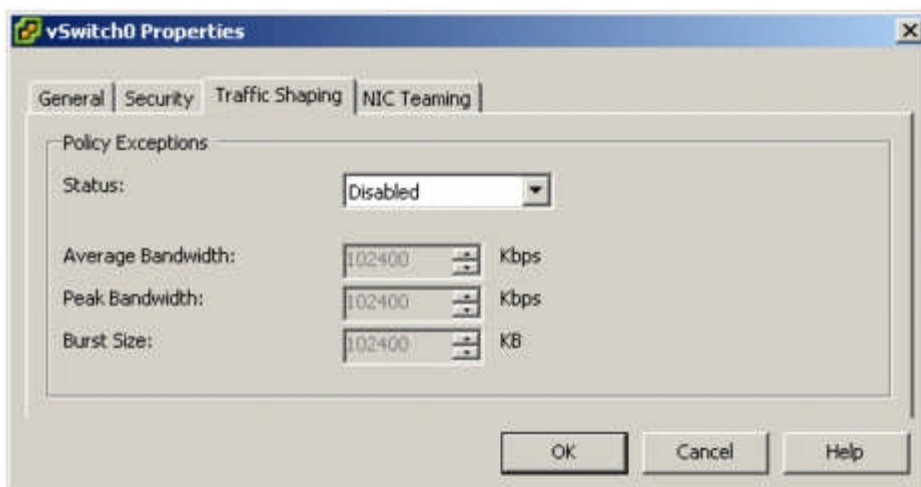
4 In the **vSwitch Properties** dialog box, click the **Ports** tab.

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

The **Properties** dialog box for the selected vSwitch appears.

6 Click the **Traffic Shaping** tab.

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



These are the policies to which the per port group exceptions are applied.

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

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

The remaining fields define network traffic parameters:

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

Load Balancing and Failover Policy

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

- Load Balancing policy

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

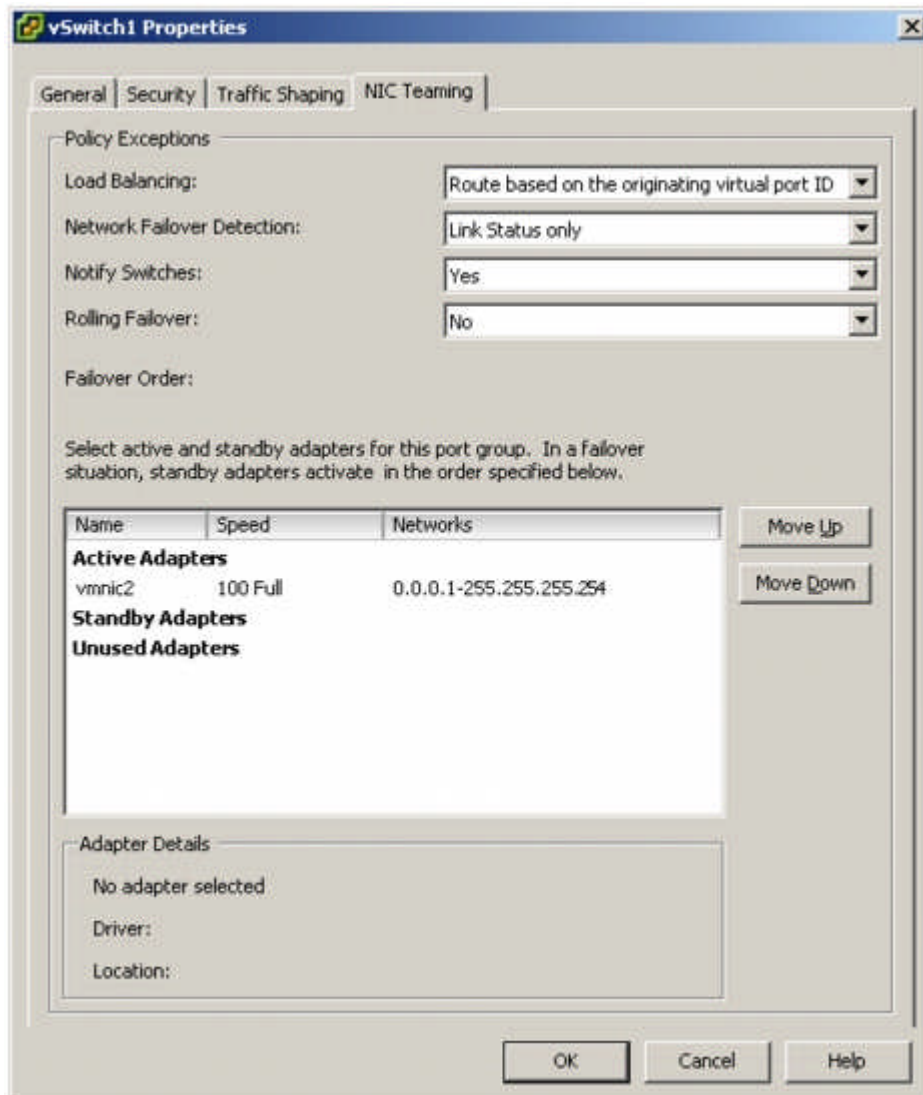
Note Incoming traffic is controlled by the Load Balancing policy on the physical switch.

- Failover Detection: Link Status/Beacon Probing
- Network Adapter Order (Active/Standby)

To edit the failover and load balancing policy

- 1 Log into the VMware VI Client and select the server from the inventory panel.
The hardware configuration page for this server appears.
- 2 Click the **Configuration** tab, and click **Networking**.
- 3 Select a vSwitch and click **Edit**.
- 4 In the **vSwitch Properties** dialog box, click the **Ports** tab.
- 5 To edit the **Failover and Load Balancing** values for the vSwitch, select the vSwitch item and click **Properties**.
The **Properties** dialog box for the vSwitch appears.
- 6 Click the **NIC Teaming** tab.

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



7 In the **Policy Exceptions** pane:

- **Load Balancing** — Specify how to choose an uplink.
 - **Route based on the originating port ID** — Choose an uplink based on the virtual port where the traffic entered the virtual switch.
 - **Route based on ip hash** — Choose an uplink based on a hash of the source and destination IP addresses of each packet. For non-IP packets, whatever is at those offsets is used to compute the hash.
 - **Route based on source MAC hash** — Choose an uplink based on a hash of the source Ethernet.
 - **Use explicit failover order** — Always use the highest order uplink from the list of Active adapters which passes failover detection criteria.
- **Network Failover Detection** — Specify the method to use for failover detection.
 - **Link Status only** — Relies solely on the link status provided by the network adapter. This detects failures, such as cable pulls and physical switch power failures, but not configuration errors, such as a physical switch port being blocked by spanning tree or misconfigured to the wrong VLAN or cable pulls on the other side of a physical switch.
 - **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 mentioned above that are not detected by link status alone.
- **Notify Switches** — Select **Yes** or **No** to notify switches in the case of failover.

If you select **Yes**, whenever a virtual NIC is connected to the vSwitch or whenever that virtual NIC's traffic would be routed over a different physical NIC in the team due to a failover event, a notification is sent out over the network to update the lookup tables on physical switches. In almost all cases, this 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.

- **Rolling Failover**— Select **Yes** or **No** to disable or enable rolling.

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

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

[VMware Infrastructure 3 Online Library](#) | [Send feedback](#) | [Technical Support](#) | Copyright © 2006–2009 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at <http://www.vmware.com/go/patents>.