

WHITE PAPER

VMware ESX Server 802.1Q - VLAN Solutions



Executive Summary

Since the release of ESX Server 2.1, ESX Server supports VLAN (IEEE 802.1Q) Trunking with ESX Server virtual switches. Using VLANs customers may enhance security and leverage their existing network infrastructures with ESX Server.

This white paper provides an overview of VLAN concepts and benefits and illustrates three possible ESX Server and virtual machine VLAN configurations. It then compares the advantages and disadvantages of the three possible configurations and recommends some best practices. The paper also includes a few configuration samples for both ESX Server and the external physical switches and finally concludes with a frequently asked questions list.

Table of Contents

VLAN Overview

ESX Server VLAN Solutions

- Virtual Machine Guest Tagging (VGT Mode)
- External Switch Tagging (EST Mode)
- ESX Server Virtual Switch Tagging (VST Mode)

VLAN Configuration

- ESX Server Configuration
 - ESX Server Configuration for VST Mode
 - ESX Server Configuration for VGT Mode
 - ESX Server Configuration for EST Mode
- Physical Switch Configuration

FAQ

VLAN Overview

VLANs provide for logical groupings of stations or switch ports, allowing communications as if all stations or ports were on the same physical LAN segment. This includes stations or ports that are physically located on different 802.1D bridged LANs.

Technically, each VLAN is simply a collision domain, also known as broadcast domain. VLAN broadcast domains are configured through software rather than hardware, so even if a machine is moved to another location, it can stay on the same VLAN broadcast domain without hardware reconfiguration. Also, traditional 802.1D bridged LANs have one single broadcast domain, so all broadcast frames are received by all stations in the network (Figure 1). VLAN networks may have multiple virtual broadcast domains within the boundary of an 802.1D bridged LAN (Figure 2).

Figure 1

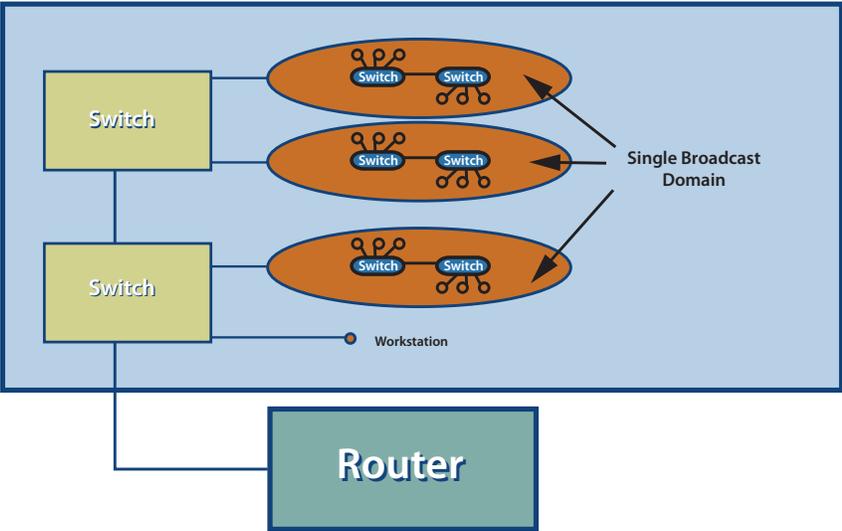
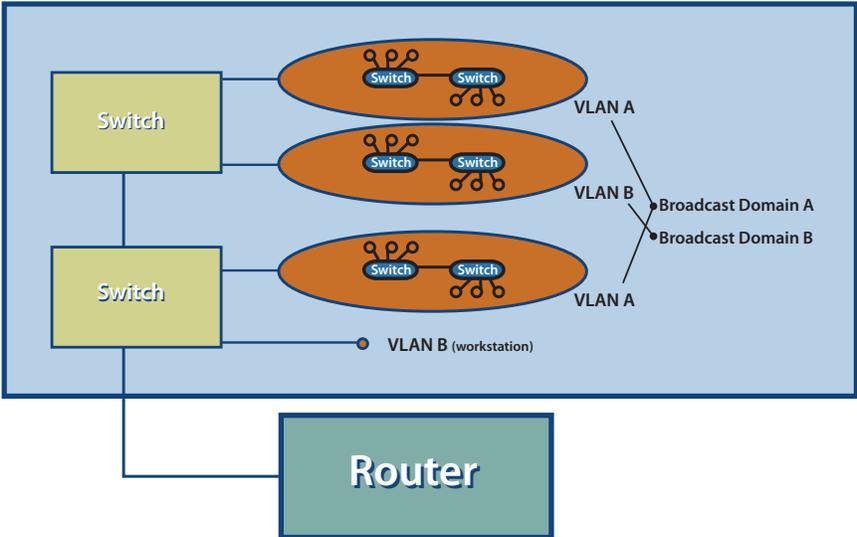


Figure 2



Major benefits of using VLANs are:

1) Flexible Network Partition and Configuration

Using VLANs, a network can be partitioned based on the logical grouping (Figure 3), not based on the physical topology. For instance, you can move a user from the sales floor to the accounting floor and maintain the same logical grouping even though the physical topology has changed.

In Figure 2 above, because none of the hosts on VLAN-A can see any traffic from VLAN-B, it is harder for any malicious users on VLAN-A to break into VLAN-B. Without VLANs, ARP spoofing and ARP poisoning is much easier.

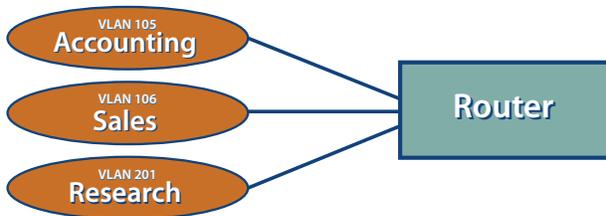


Figure 3

2) Performance Improvement

TCP/IP network protocols and most other protocols broadcast frames periodically to advertise or discover network resources. On a traditional flat network, frames reach all hosts on a network. This can have a significant impact on the network performance with a large number of end users. Confining broadcast traffic to a subset of the switch ports or end users saves lots of network bandwidth and processor time.

3) Cost Savings

Without VLANs, network administrators partition LANs into multiple broadcast domains by using routers between those segments. However, routers are expensive and may introduce more delay.

Some early proprietary VLAN implementations were restricted to a single switch and tagging packets based on physical ports. IEEE 802.1Q tagging can span VLAN across switches or even across WANs.

In order to extend VLANs across switches, a trunk link must interconnect the switches. Frames on the trunk are encapsulated in the IEEE 802.1Q format and are not much different from the regular Ethernet frames except that they contain an extra four bytes inserted after the source and destination MAC address (Figure 4). In the four byte 802.1Q tag, the first two bytes (0x8100) are an indicator that the following frame is an 802.1Q frame and the next two bytes are for the VLAN tag (3 bits for priority bits, 1 bit for Canonical Format Indicator, and last 12 bits for the VLAN ID). VLAN ID 0 is reserved for user priority

data, which is not supported by ESX Server; VLAN ID 4095 is reserved for future definition. The special native VLAN issue is discussed separately later.

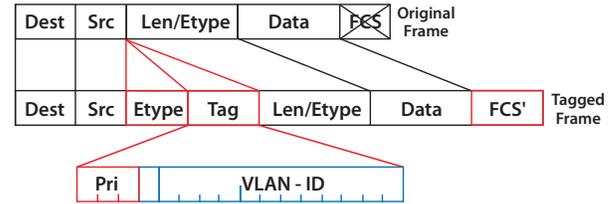


Figure 4

Currently there are many forms of VLAN tagging, and they can be categorized based on the tagging algorithm:

1) Port-Based VLAN

For example, one could group switch ports 1 to 2 into a VLAN 101 for the HR department and ports 6 to 12 into VLAN 102 for the IT department. This kind of configuration is the simplest to deploy and maintain. However, it is inflexible as moving the workstation may require changing the switch configuration.

2) MAC-Based VLAN

Tagging based on Layer 2 MAC address. This requires significant initial configuration of the switches,. However, automatic tracking is possible thereafter.

3) Protocol Based VLAN

Tagging based on layer 3 IP address, layer 4 transport protocol information, or even higher layer protocol information.

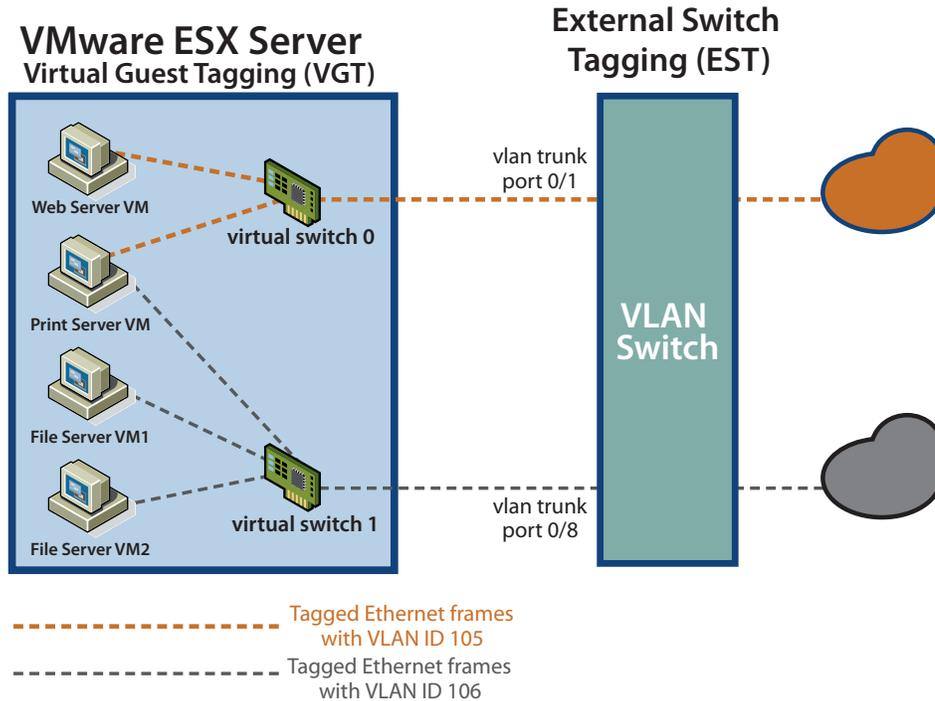
4) Policy Based VLAN

Tagging based on certain policies or user configuration. This may involve classifying network traffic into groups and assigning QoS priority bits and VLAN ID to each group.

ESX Server VLAN Solutions

In order to support VLANs for ESX Server users, one of the elements on the virtual or physical network has to tag the Ethernet frames with an 802.1Q tag. There are three different configuration modes to tag (and untag) the packets for virtual machine frames.

Figure 5



Virtual Machine Guest Tagging (VGT Mode)

The user may install an 802.1Q VLAN trunking driver inside the virtual machine. This preserves tags between the virtual machine networking stack and the external switch when frames are passed from and to virtual switches. (Figure 5)

The advantages of using this mechanism are:

1. The number of VLANs per virtual machine is not limited to the number of virtual adapters, which means your virtual machines can be on any number of VLANs on your network.
2. If a physical server is already running the VLAN driver, then it is natural to use P2V Assistant to convert this server and keep running the existing VLAN tagging.

The disadvantage of using this mechanism is that it is not always possible or easy for the user to find and configure 802.1Q drivers for the guest operating system.

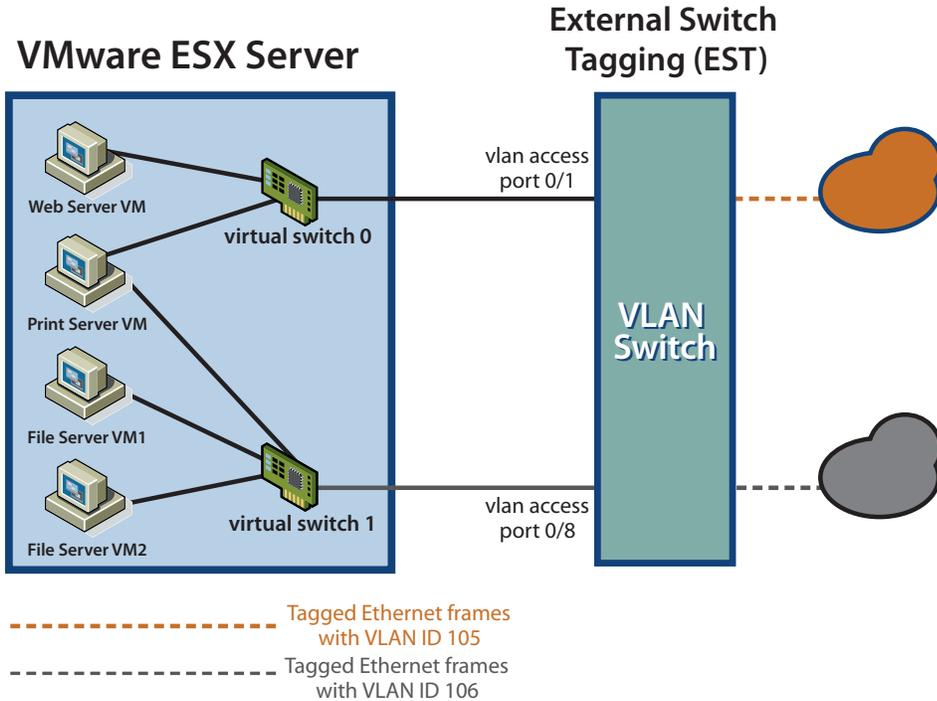
Without VLAN hardware acceleration, it takes significant CPU cycles to tag the outbound frames and remove the tag for inbound frames.

Customers may have to use this solution if a single virtual machine must be on five or more different VLANs in the network. This method could also be appropriate when physical servers running 802.1Q trunking drivers exist already are being virtualized. Such servers can be virtualized simply by using P2V Assistant and no additional network configuration is required. The new virtual machine will automatically inherit all VLAN settings from the physical machine.

Some operating systems, including some Linux distributions, support 802.1Q trunking well.

Note: VGT mode is supported in ESX Server 2.1.1 or later releases, but not in ESX Server 2.1.0 or earlier releases.

Figure 6



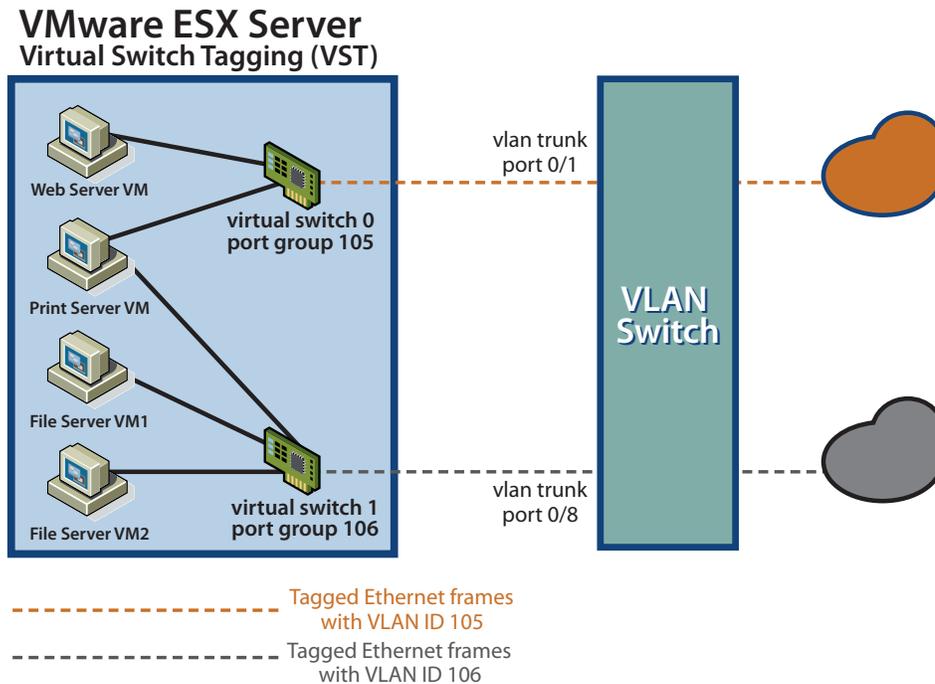
External Switch Tagging (EST Mode)

The user may use the external switches for VLAN tagging. This is similar to a physical network, and VLAN configuration is normally transparent to each individual physical server. The tag is appended when a packet arrives at a switch port and stripped away when a packet leaves a switch port toward the server. (Figure 6)

ESX Server users can set up the VLAN configuration as they would for any physical server. This method was recommended before VLAN support was added to virtual switches in ESX Server. One drawback of this approach is that if port-based VLAN tagging is used (common in the enterprise VLAN deployment), the total number of virtual LANs supported would be limited to the number of NICs installed on a given ESX Server system. This limitation is removed using VGT or VST modes as described in this white paper.

Note: EST mode is supported in all ESX Server releases.

Figure 7



ESX Server Virtual Switch Tagging (VST Mode)

In this mode, the user provisions one port group on a virtual switch for each VLAN and then attaches the virtual machine's virtual adapter to the port group instead of the virtual switch directly. The virtual switch's port group tags all outbound frames and removes tags for all inbound frames. It also ensures that broadcast frames on one VLAN do not leak into a different VLAN. (Figure 7)

ESX Server virtual switch tagging has the following benefits:

1. Different VLAN frames can be multiplexed onto a single physical NIC, so you can consolidate all traffic regardless of VLAN into a single physical NIC. There is no need for multiple NICs to service multiple VLANs.
2. It eliminates the need to run a guest operating system specific VLAN driver inside the virtual machine.
3. Since all modern high-speed network cards support VLAN acceleration, there is little performance impact by supporting VST mode in the virtual switches.
4. Once the switch trunk mode is appropriately set up, no additional switch configuration is needed when provisioning additional VLANs. External switch configuration becomes easy.
5. VST mode does not compromise networking or overall system performance, provided you use NICs that support VLAN hardware acceleration.

Note: VST mode is supported in ESX Server 2.1.0 and later releases. Even though ESX Server 2.1.1 and later releases support all three solutions illustrated in this section (VST, EST and VGT modes), in general, VST mode is the only VLAN solution recommended by VMware. However, due to the legacy configuration or existing network configuration constraints, VST mode may not be possible or the most appropriate solution in all environments.

VLAN Configuration

ESX Server Configuration

ESX Server Configuration for VST Mode

Virtual Switch Tagging (VST) and Virtual Machine Guest Tagging (VGT) are mutually exclusive features. By default VST is enabled.

ESX Server 2.1 introduced the ability to configure port groups in the VMware Management Interface. Port groups assist in VLAN configuration. Users can create a new port group in the Network Connections page on the Virtual Switches tab. The Configuring a port group involves specifying a Port Group Label and VLAN ID (Figure 8). These values must be unique on a virtual switch. Once a port group is created, users can use the VLAN network label (that is, the port group label) when configuring a virtual machine.

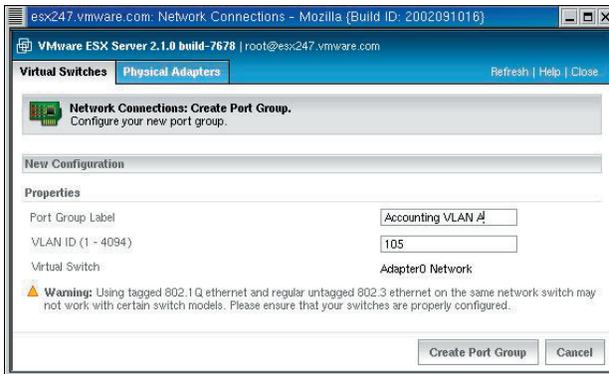


Figure 8

In ESX Server 2.x, the VLAN ID is the only characteristic associated with a port group, so the port group serves nothing but the VLAN. That is, you do not need to create a port group unless you want to use VST mode. In the future, however, there will be more policies associated with a port group, so you may need to create port groups even if you do not use VLAN.

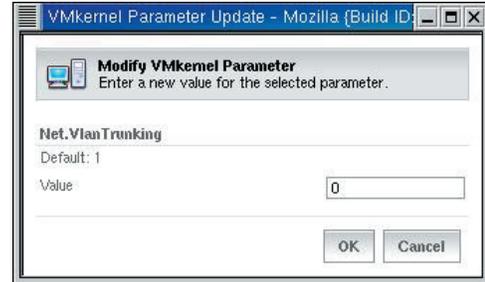


Figure 9

ESX Server Configuration for VGT Mode

VGT mode is disabled by default. In order to use VGT mode, you must disable the `Net.VlanTrunking` setting in the Advanced Settings page of the VMware Management Interface (see Figure 9). After rebooting the ESX Server system, you can then run the 802.1Q VLAN trunking driver inside the virtual machine.

This is not a recommended configuration for most customers. Please review the description of this mode in the previous section of this white paper. It should be used only in special circumstances as described in that section.

ESX Server Configuration for EST Mode

There is nothing to configure in ESX Server in order to use External Switch Tagging (EST mode).

For example, for port-based EST mode, you may simply allocate and connect one NIC port to one switch VLAN port. Since it is a one to one relationship, the number of VLANs supported on your ESX Server system is limited to the number of physical NIC ports assigned to the `vmkernel`.

Physical Switch Configuration

This section explains the external physical switch configuration for VST mode only. Configuration for VGT mode may differ slightly. The command line syntax for different switches varies too.

The link between a physical NIC on an ESX Server system and an external switch port is considered an inter-switch link.

VLAN configuration is different for switches from different vendors or of different types. For purposes of illustration in this white paper, we provide a few sample configuration snippets from Cisco switches (CatOS or IOS), but the actual syntax could be different from your own switch syntax. Please refer to your switch manual for more information.

1) Specifying Trunk Ports

In order to interoperate with Virtual Switch Tagging (VST mode), you must configure the external switch ports to be VLAN Trunk ports. VST mode does not support Dynamic Trunking Protocol (DTP), so you have to make the trunk statically and unconditionally.

In the following example on a Cisco switch running CatOS, only ports 0/1 and 0/2 are good for VST mode, and ports 1/3, 1/4, and 1/5 do not pass VLAN frames properly to the ESX Server systems.

```
CatOS Console> (enable) set trunk
0/1 nonegotiate dot1q
CatOS Console> (enable) set trunk
0/2 on dot1q
CatOS Console> (enable) set trunk
0/3 off dot1q
CatOS Console> (enable) set trunk
0/4 desirable dot1q
CatOS Console> (enable) set trunk
0/5 auto dot1q
```

The *desirable* and *auto* do not work because the switch expects its peer (that is, the ESX Server virtual switch port) to communicate by way of DTP. The *nonegotiate* and *on* options enable VLAN trunking unconditionally. The difference between the *nonegotiate* and *on* options is that *on* mode still sends out DTP frames. To minimize the unnecessary network traffic, use the *nonegotiate* option.

Please note that you may not always find all the options mentioned above. For example, *dot1q* may be the only protocol supported on your switches, so you do not need to specify *dot1q*.

2) Specifying VLANs for Trunking

When you put a port into trunk mode, you must make sure that the VLANs you have configured on your ESX Server system are defined and allowed by the switch trunk port. The default behavior varies among different types of switches and between vendors.

You may need to define all the VLANs used with ESX Server on the physical switch explicitly. For each VLAN definition, you may specify the VLAN ID, name, type, MTU, security association identifier (SAID), state, ring number, bridge identification number, and so on.

For example:

```
CatOS Console> (enable) set vlan
105 name accounting type ethernet
mtu 1500 said 100105 state inactive
```

For switches that allow all ports by default (for example, VLAN 1 – VLAN 1005 on some Cisco switches), you may not need to do anything. However for best security practice, VMware recommends you restrict the VLANs only to those you need.

The following example shows you how to clear all VLANs first and then enable VLANs 80, 81, 82, 83, 84, 85, 105, 106 and 303 on port 1/1:

Step 1. Clearing all VLANs allowed:

```
IOS Console> (enable) switchport
trunk allowed remove 1-1005
CatOS Console> (enable) clear
trunk 1/1 1-1005
```

Step 2. Adding VLANs desired:

```
IOS Console> (enable) switchport
trunk allowed add 10-15,20,25
CatOS Console> (enable) set trunk
1/1 on dot1q 80-85,105-106,303
```

For those switches that none of the VLAN IDs are allowed by default, you have to add the VLANs to the trunk explicitly, so step 1 above is not needed.

3) Native VLAN Issue (a.k.a. "VLAN 1 Issue")

Native VLAN is used for switch control and management protocol. Native VLAN frames are not tagged with VLAN IDs in many types of switches, and in which case the trunk ports implicitly treat all untagged frames as the native VLAN frame.

VLAN 1 is the default native VLAN ID for most Cisco switches. However, in many enterprise networks, the native VLAN is VLAN 1 or 100, it could be any number depending on your switch type and running configuration.

It is a common best practice to avoid using native VLAN (often VLAN 1) for any regular data traffic. VMware recommends you not associate any ESX Server virtual switch port group VLAN IDs with the native VLAN. Also, as long as you avoid using native VLAN for your VLAN port groups, there is no native VLAN related configuration necessary on the ESX Server systems.

In the event that you have to associate VLAN 1 with a port group and pass virtual machine network traffic through it, you must do either one of the following two things:

- Make sure VLAN 1 is not the native VLAN on your physical switches. You may change the default native VLAN to another VLAN ID. For example, to change the native VLAN ID to 101, use the following command:

```
IOS Console> (enable) switchport
trunk native 101
```

- Enable the native VLAN 802.1q tagging capability. Some switches do not support this option and some other switches do not need it as tagging on native VLAN is enabled by default.

```
IOS Console> (enable) vlan dot1q
tag native
CatOS Console> (enable) set dot1q-
all-tagged enable
```

Note that when you change the behavior of the native VLAN on one of your external switches by doing either step above, it is likely that you must change all the neighbor switches as well so that they can still communicate on the native VLAN properly.

FAQ

Unless specified otherwise, answers are for ESX Server Virtual Switch Tagging (VST mode) only.

Q: If I do not use VLANs on my ESX Server system, what do I need to do? Do I need to disable `Net.VlanTrunking` in the VMware Management Interface?

A: There is nothing you need to do if you do not use VLANs with ESX Server. The `Net.VlanTrunking` settings in the Advanced Settings page (on the Options tab) needs to be disabled only if you decide to use Virtual Machine Guest Tagging (VGT mode). If you disable the setting, you must reboot the ESX Server system.

Q: What NIC controllers are supported for the VLAN modes and which ones support hardware acceleration?

A: All of the Intel, Broadcom, and 3Com NIC controllers support External Switch Tagging (EST mode) in all ESX Server releases. All of the Intel and Broadcom NIC controllers support both Virtual Switch Tagging (VST mode) and Virtual Machine Guest Tagging (VGT mode) in ESX Server 2.1.1 and later. Also, all the Gigabit NICs listed in the ESX Server I/O Compatibility Guide support VLAN hardware acceleration in VST mode.

Q: Can the ESX Server VLAN modes work with NIC teaming?

A: All three VLAN modes work with NIC teaming seamlessly. In VST mode, the teamed virtual switch uplinks do not create loops, so it is best to disable Spanning Tree Protocol (or enable PortFast) on the external switch ports that are connected to the ESX Server system.

Q: Can a virtual machine be configured on multiple VLANs?

A: You can only configure one VLAN ID for each virtual machine's virtual network adapter. However, since you can configure up to four virtual adapters per virtual machine, you can set up a virtual machine spanning four different VLANs.

Q: What is the valid VLAN ID range supported on an ESX Server system?

A: The VLAN ID range defined in the IEEE 802.1Q specification is from 1 to 4094. VST mode supports VLAN ID ranges from 1 to 4094. In practice, this range is larger than most switches can handle, so make sure your switch can accept the VLANs you configure on your ESX Server systems. Be especially careful

about using the native VLAN, which may require special switch configuration support. For best results, you may want to avoid using native VLAN for regular virtual machine data traffic.

Q: How many different VLAN IDs can an ESX Server virtual switch support? That is, how many port groups can I add to a virtual switch?

A: Port groups can be on any valid IEEE 802.1Q VLAN ID (1 – 4094), so you can add up to 4094 port groups to any virtual switch. However since there are only up to 32 active ports or virtual adapters per virtual switch, it is not always true that the more port groups per virtual switch you have configured statically, the more port groups you can use simultaneously.

Q: Is there any performance penalty caused by running VLAN trunking in ESX Server?

A: No. There is no measurable performance impact for using VST mode.

Q: Are 802.1Q priority bits supported in ESX Server?

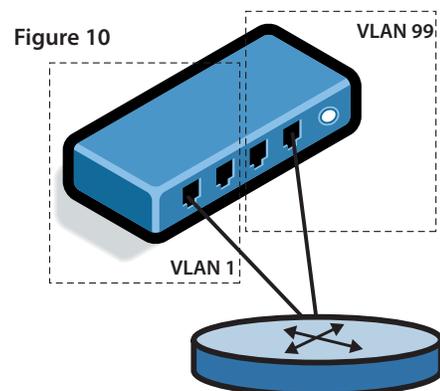
A: No.

Q: Can I hot swap the VLAN ID of a virtual adapter when a virtual machine is running?

A: Yes. Because you can hot swap the network port group label of a virtual machine dynamically, you can change VLAN attributes such as the VLAN ID, dynamically.

Q: How can I set up communication between VLANs?

A: A router (Figure 10), a layer 3 switch, or a switch that supports communication between VLANs must be involved. The solution is typically proprietary and vendor-specific.



Q: What is the difference between VLANs and IP subnets?

A: VLAN is a layer 2 technology only and IP subnets operate at layer 3. They are orthogonal. However, it is common to have a one to one relationship between a VLAN and an IP subnet though this is not required.

Q: Does ESX Server support VLANs for non-IP networks?

A: Yes. VLAN is a layer 2 technology that works with any layer 3 networks. You may provision virtual machines without TCP/IP stacks, such as NetWare, onto VLANs.

Q: Can I configure my virtual adapter to be on multiple VLANs?

A: Not if you are using ESX Server Virtual Switch Tagging (VST mode). But for Virtual Machine Guest Tagging (VGT mode), you may configure multiple VLANs for one virtual adapter inside the virtual machine.

Q: How can I configure my virtual machine to be on multiple VLANs?

A: If you use VST mode, you may add up to four virtual adapters in a virtual machine with each of them on a different VLAN port group. If you want to configure your virtual machines to be on more than four VLANs, you must use VGT mode. As described above, you cannot use VGT mode for some virtual machines and VST mode for the other virtual machines on the same ESX Server system.

Q: What is the `Net.SwitchFailoverBeaconVlanID` option in the Advanced Settings page (on the Options tab) in the VMware Management Interface?

A: This option is used very rarely. In ESX Server 2.1 and later, if you use NIC teaming and enable beacon monitoring protocol (disabled by default), you may then need to set `Net.SwitchFailoverBeaconVlanID` to one of the VLAN IDs your external switches allow. Otherwise, you may ignore this option.

Q: Can I migrate a virtual machine with VMotion if the virtual machine's virtual network adapters use VLANs (port groups)?

A: Yes. Make sure that the destination ESX Server system has the same port groups defined and that the external switch is correctly configured for VLANs.

Q: Can I send network traffic from migrating a virtual machine with VMotion over VLANs?

A: Yes. You can send such traffic over any virtual switch port groups you have defined. For best security, VMware recommends you use a dedicated virtual switch or, at minimum, a dedicated virtual switch VLAN port group for VMotion.

Q: All the VLAN port groups work for me except one VLAN ID. Why?

A: It is likely that the VLAN that does not work for you is the native VLAN in your network. See Native VLAN Issue on page 8 for more information.

Q: Is Cisco's Inter-Switch Link (ISL) Protocol supported by ESX Server virtual switches?

A: No.

Q: Is Dynamic Trunking Protocol supported by ESX Server virtual switches?

A: No, VMware does not support it for network security and stability reasons.

Q: Do any ESX Server virtual switches support per-VLAN Spanning Tree Protocol (STP)?

A: No, ESX Server virtual switches do not support Spanning Tree Protocol. Multiple virtual switches on a single ESX Server system do not create any loops when they connect to the external VLAN or non-VLAN networks.

Q: How can I provision VLANs in the service console?

A: Normally you use External Switch Tagging (EST mode) in the service console. However, if your service console network interfaces are created by the `vmxnet_console` driver, you can use Virtual Switch Tagging (VST mode).

Q: Do I have to connect a trunk port to an ESX Server system?

A: Yes, if you want to power on a virtual machine using port groups (that is, in VST mode). Some users have security concerns when connecting trunk ports to servers. However, ESX Server virtualizes both servers and switches, so the link between an ESX Server virtual network adapter and a switch port is considered an interswitch link.

Q: Can I add the same VLAN ID to multiple virtual switches on the same ESX Server system?

A: Yes.

Q: Can I provision an 802.1Q trunk directly between two virtual switches on the same ESX Server system?

A: No. Because none of the virtual switches on an ESX Server system are connected, there is no way to provision any 802.1Q trunks among them directly. For the same reason, ESX Server virtual switches are loop-free.



VMware, Inc. 3145 Porter Drive Palo Alto CA 94304 USA Tel 650-475-5000 Fax 650-475-5001 www.vmware.com
Copyright © 2004 VMware, Inc. All rights reserved. Protected by one or more of U.S. Patent Nos. 6,397,242 and 6,496,847; patents pending. VMware, the VMware "boxes" logo, GSX Server and ESX Server are trademarks of VMware, Inc. Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation. Linux is a registered trademark of Linus Torvalds. All other marks and names mentioned herein may be trademarks of their respective companies.

