WHITE PAPER

# VMware ESX Server 3 802.1Q VLAN Solutions

vmware®

# Executive Summary

The virtual switches in ESX Server 3 support VLAN (IEEE 802.1Q) trunking. Using VLANs, you can enhance security and leverage your 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 configuration samples for both ESX Server and the external physical switches and concludes with a list of frequently asked questions.

## Table of Contents

## 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 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).

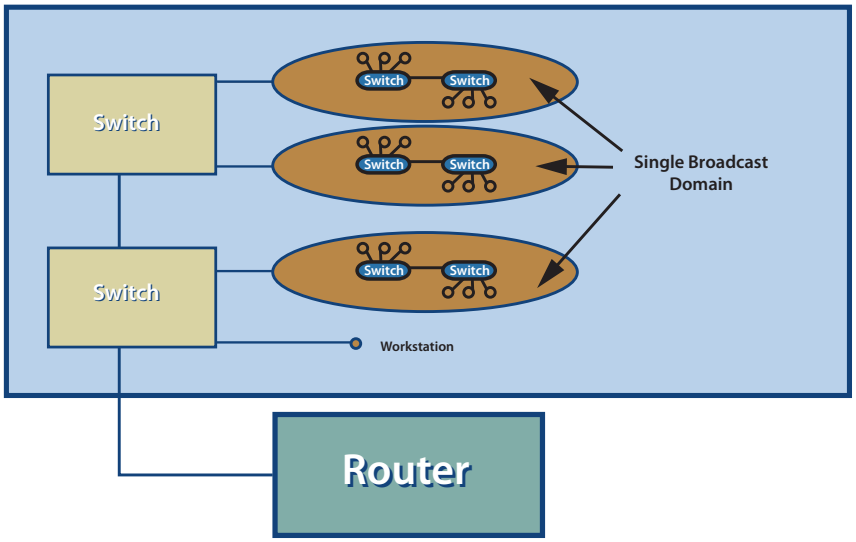


**Figure 1**

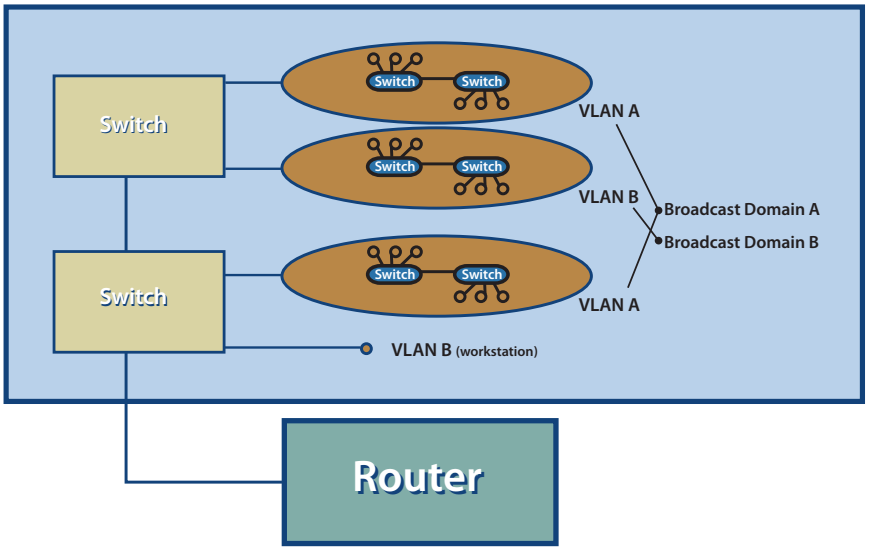


**Figure 2**

Major benefits of using VLANs are:

- **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 are much easier.

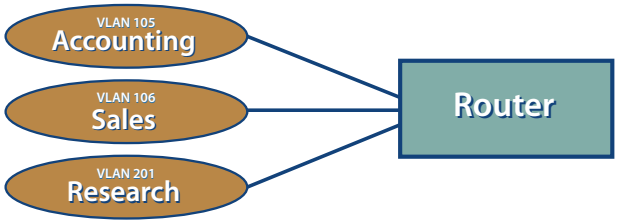


**Figure 3**

- **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 when you have a large number of end users. Confining broadcast traffic to a subset of the switch ports or end users saves significant amounts of network bandwidth and processor time.
- **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 a 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 (three bits for priority bits, one 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.

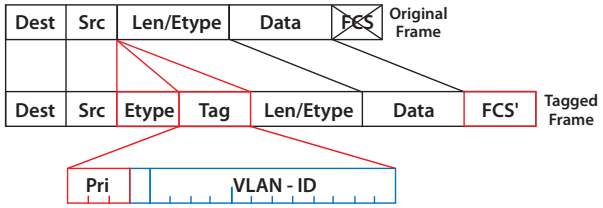


**Figure 4**

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

- **Port-Based VLAN**
  Tagging based on switch ports. For example, you might group switch ports 1 to 2 into 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 a workstation may require changing the switch configuration.
- **MAC-Based VLAN**
  Tagging based on layer 2 MAC address. This requires significant initial configuration of the switches,. However, automatic tracking is possible thereafter.
- **Protocol Based VLAN**
  Tagging based on layer 3 IP address, layer 4 transport protocol information, or even higher-layer protocol information.
- **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 802.1Q tag. There are three different configuration modes to tag (and untag) the packets for virtual machine frames.

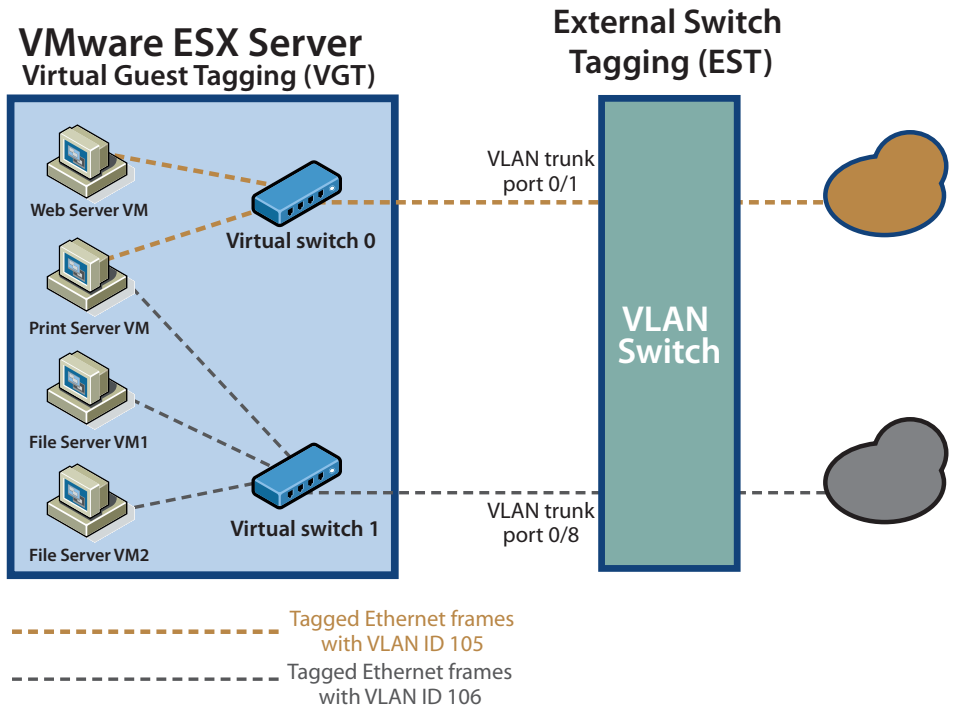


**Figure 5**

### *Virtual Machine Guest Tagging (VGT Mode)*

You may install an 802.1Q VLAN trunking driver inside the virtual machine, and tags will be preserved between the virtual machine networking stack and external switch when frames are passed from/to virtual switches. (Figure 5)

The advantages of using this mechanism are:

- 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.
- If a physical server is already running VLAN driver, then it will be 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 additional CPU cycles to tag the outbound frames and remove the tag for inbound frames.

You 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 existing physical servers running 802.1Q trunking drivers 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.

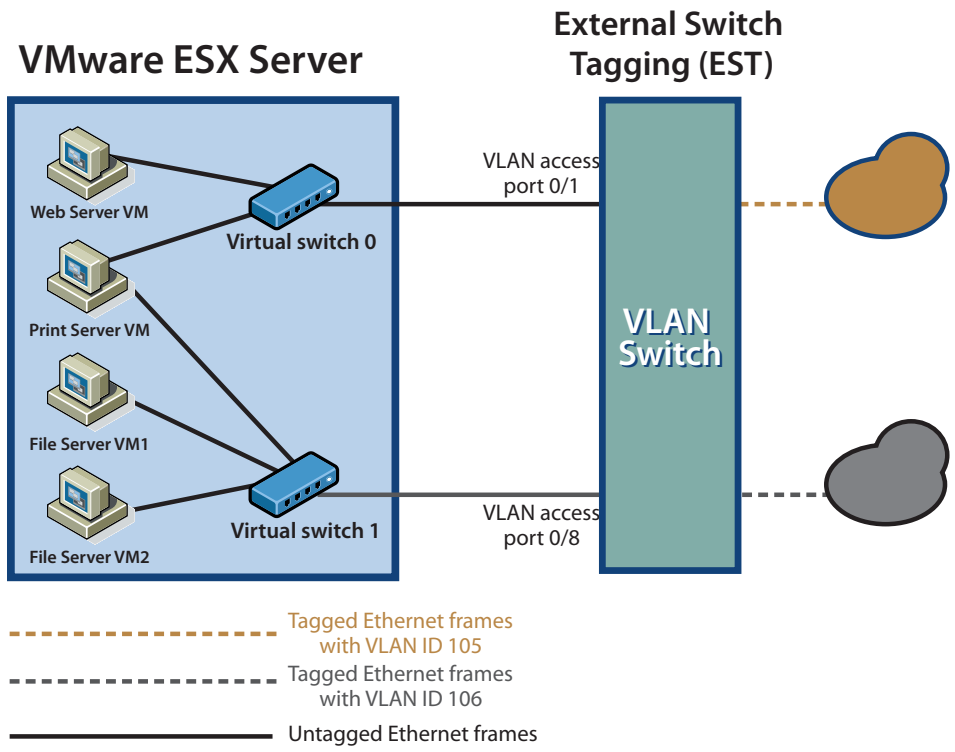


**Figure 6**

### *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 just as they would for any physical server. One drawback of this approach is that if port-based VLAN tagging is used (common in enterprise VLAN deployments) 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.

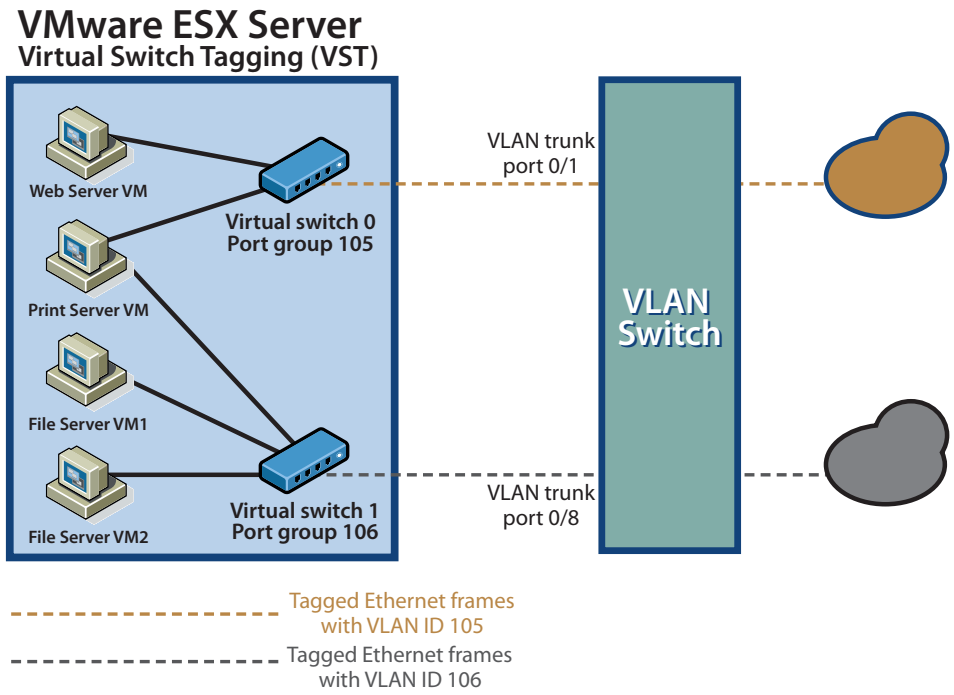


**Figure 7**

### *ESX Virtual Switch Tagging (VST Mode)*

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

ESX Server virtual switch tagging has the following benefits:

- 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 VLAN's.
- It eliminates the need to run a guest operating system specific VLAN driver inside the virtual machine
- Since all modern high-speed network cards support VLAN acceleration, there is little performance impact by supporting VLAN tagging in the virtual switches.
- Once the switch trunk mode is appropriately set up, no additional switch configuration is needed when provisioning additional VLANs. External switch configuration becomes easy.

# VLAN Configuration

## ESX Server Configuration

### *ESX Server Configuration for VST*

Virtual Switch Tagging is enabled when the port group's VLAN ID is set to any number between 1 and 4094, inclusive.

To use VST, you must create appropriate port groups. You can give each one a port group label and a VLAN ID. The port group values must be unique on a virtual switch. Once you have created a port group, you can use the port group label in the virtual machine configuration.

To set port group properties

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, then click **Networking**.

3) On the right side of the window, click **Properties for a network**.

   The vSwitch Properties dialog box appears.

4) Click the **Ports** tab.

5) Select the port group and click **Edit**.

6) In the Properties dialog box for the port group, click the **General** tab to change:

   - **Network Label** — a name that identifies the port group that you are creating.
   - **VLAN ID** — identifies the VLAN that the port group's network traffic will use.

7) Click **OK** to exit the vSwitch Properties dialog box.

### *ESX Server Configuration for VGT*

To use VGT, enter 4095 as the port group's VLAN ID. 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*

You do not need to configure anything in ESX Server in order to use External Switch Tagging. It is enabled by default.

For example, for Port-based EST, 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.

EST is enabled when the port group's VLAN ID is set to 0 or left blank.

## Physical Switch Configuration

This section explains the external physical switch configuration for Virtual Switch Tagging (VST) only. Configuration for Virtual Machine Guest Tagging (VGT) may differ slightly. The command line syntax for different switches varies too.

The link between an ESX physical NIC and an external switch port is considered an interswitch link.

VLAN configuration is different for switches of different vendor or type. For illustration purpose, we will give you 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.

**Specifying Trunk Ports**

In order to interoperate with 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 static and unconditional.

In the following example on a Cisco switch running CatOS, only ports 0/1 and 0/2 are good for VST mode, and ports 0/3, 0/4, and 0/5 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` settings do not work because the switch expects its peer (that is, the ESX Server virtual switch port) to communicate using DTP. The `nonegotiate` and `on` options enable VLAN trunking unconditionally. The difference between `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 instance, `dot1q` may be the only protocol supported on your switches, so you do not need to specify `dot1q`..

**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 explicitly on the physical switch. 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 to only those you need.

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

1) Clear all VLANs allowed:

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

2) Add 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 switches on which 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.

**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 any VLAN ID in many types of switches. In these cases the trunk ports implicitly treat all untagged frames as native VLAN frames

VLAN 1 is the default native VLAN ID for most Cisco switches. However, in many enterprise networks, the native VLAN might be 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 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 the 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, you will likely need to change all the neighbor switches as well so they can still communicate on the native VLAN properly.

# FAQ

Unless specified otherwise, answers are for VST mode only.

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

A: All of the Intel, Broadcom, and 3Com NIC controllers support EST mode. All of the Intel and Broadcom NIC controllers support both VST mode and VGT mode. Also, all the Gigabit Ethernet NICs on the ESX Server I/O Compatibility Guide support VLAN hardware acceleration in VST mode.

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

A: All three VLAN modes work seamlessly with NIC teaming. 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 configure only one VLAN ID for each virtual network adapter on a virtual machine. 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 the native VLAN for regular virtual machine data traffic.

Q: How many different VLAN IDs can an ESX Server virtual switch support?

A: ESX Server 3 supports 4094 VLAN IDs on each virtual switch.

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. In the port group settings, you can change the VLAN ID while a virtual machine is running.

Q: How can I set up communication between VLANs?

A: A router (Figure 8), a layer 3 switch, or a switch that supports communication between VLANs must be involved. One approach, for example, is to use a software router running in a virtual machine on an ESX Server host.

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

A: VLAN is a layer 2 technology only, and IP subnets operate at

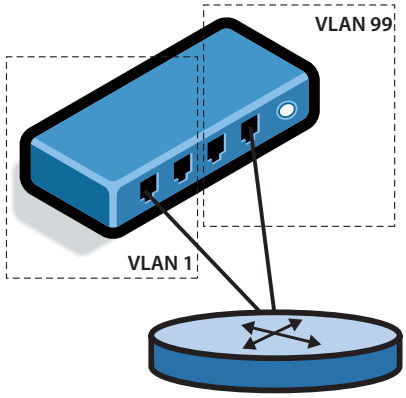


**Figure 8**

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 — for example, virtual machines running NetWare — onto VLANs.

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

A: Not if you are using VST mode. But for VGT mode, you can 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 one 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.

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 used to migrate 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 a 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, we do not support it for network security and stability reason.

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 to non-VLAN networks.

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

A: Settings for the service console are handled the same way that they are for virtual machines. Set the appropriate value for the VLAN ID in the port group settings for the service console.

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 VST or VGT 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 same VLAN ID to multiple virtual switches of the same ESX Server system?

A. Yes.

Q. Can I provision a 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.

VI3-ENG-Q206-234

**VMware, Inc. 3145 Porter Drive Palo Alto CA 94304 USA Tel 650-475-5000 Fax 650-475-5001 www.vmware.com**



vmware®