

Planning and Implementing VLANs with HP-UX

HP Part Number: 5992-0538
Published: March 2007



Table of Contents

About This Document.....	5
What is VLAN?.....	7
Benefits of VLANs	8
VLAN-Aware Switches Are the Key.....	8
Which VLAN Does a Frame Belong To?.....	9
How Does a VLAN-aware Switch Work?.....	9
VLAN Tagging.....	10
Standards and Interoperability.....	10
VLAN Trunking.....	10
VLANs on HP-UX	11
Features and Advantages	11
Implementing VLANs on HP-UX.....	11
Common Usage Scenarios.....	11
Determining a Usage Scenario that Meets Your Network Needs.....	12
Priority and Class of Service (CoS).....	12
IP ToS and 802.1p Conversion—End-to-End Class of Service.....	13
Typical Customer Configurations.....	13
A VLAN Implementation Example.....	14
Using HP-UX VLANs with HP Auto Port Aggregation (APA).....	15
Using HP-UX VLANs with HP Virtual Machines (HPVM).....	16
Future HP-UX VLAN Feature Additions.....	18

About This Document

This white paper presents network managers with an overview of HP-UX Virtual LAN (VLAN) software for HP servers. The following HP-UX VLAN topics are addressed:

- An introduction to VLAN technology and its benefits—since a VLAN-aware switch is the building block of a VLAN network environment, this white paper explains how switches implement VLANs.
- HP-UX VLANs: features and advantages and how HP-UX fits into a VLAN environment, including:
 - Using HP-UX VLANs with HP Auto Port Aggregation (APA)
 - Using HP-UX VLANs with HP Virtual Machines (HPVM)
- Planning a VLAN with HP-UX servers.

Intended Audience

This document is intended for system and network administrators responsible for configuring and administering HP-UX VLANs. Administrators are expected to have knowledge of HP-UX and networking concepts, commands and configuration.

This document is not a tutorial.

Related Information

The following documents and web sites contain useful and related information:

- <http://www.hp.com/go/vlan>
- [HP-UX VLAN Documentation](#)
- [HP Auto Port Aggregation Documentation](#)
- [HP Integrity Virtual Machines Documentation](#)

What is VLAN?

Virtual LAN (VLAN) technology allows network administrators to separate logical network connectivity from physical connectivity. This concept is different from a traditional LAN in that a LAN is limited by its physical connectivity. All users in a LAN belong to a single broadcast domain¹ and can communicate with each other at the Data Link Layer or “Layer 2”. Network managers have used LANs to segment a complex network into smaller units for better manageability, improved performance, and security. For example network managers use one LAN for each IP subnet in their network. Communication between subnets is made possible at the Network Layer or “Layer 3”, using IP routers.

A VLAN can be thought of as a single physical network that can be logically divided into discrete LANs that can operate independently of each other.

Figure 1 Using VLANs to Create Independent Broadcast Domains Across Switches




Figure 1 highlights several key differences between traditional LANs and VLANs.

- All switches are interconnected to each other. However, there are three different VLANs or broadcast domains on the network. Physical isolation is not required to define broadcast domains. If Figure 1 was a traditional LAN without VLAN-aware switches, all stations would belong to one broadcast domain.
- All switch ports can communicate with one another at the Data Link Layer, if they become members of the same VLAN.
- The physical location of an end station does not define its LAN boundary.
 - An end station can be physically moved from one switch port to another without losing its “view of the network”. That is, the set of stations it can communicate with at the Data Link Layer remains the same, provided that its VLAN membership is also migrated from port to port.
 - By reconfiguring the VLAN membership of the switch port an end station is attached to, you can change the network view of the end station easily, without requiring a physical move from port to port.

1. A LAN is a broadcast domain at the Data Link Layer because a broadcast or multicast frame sent from a station is seen by all other stations in its LAN.

Benefits of VLANs

The key benefits of using VLANs include the following:

- **Bandwidth preservation:** A well-designed VLAN helps restrict broadcast and multicast traffic to only those stations listening to and responding to the traffic related to that VLAN. The network and computing resources of nonparticipating stations are unaffected, thus improving performance.
- **Manageability:** Moves, additions, and changes to network topology do not require physical changes to network topology. User mobility is much simpler because of the dynamic nature of VLANs.

Physically dispersed work groups can be logically connected within the same broadcast domain to appear as if they are on the same physical LAN. A single physical link can simultaneously serve several IP subnets when subnet-based VLANs are configured on that link. End stations using VLANs can offer rudimentary Class of Service (CoS) locally by prioritizing traffic for certain activities.

- **Enhanced security:** You can construct different security domains to provide various levels of security in the network, because the network design is more flexible than that of traditional LANs. Since frames are passed to a destination port only if the port belongs to the same VLAN as the frame, VLANs help enforce traffic isolation, thereby providing an added level of security in the network.

VLAN-Aware Switches Are the Key

To implement a VLAN in your network, you must use VLAN-aware switches. This section describes how VLAN-aware switches are different from traditional switches.

To understand how logical partitioning of a LAN infrastructure is done using VLAN, it is helpful to remember the fundamental operation of a traditional switched LAN. Without going into the details of switch design, the two rules to remember regarding the functioning of a regular LAN switch are:

1. When the switch receives a broadcast or multicast frame from a port, it floods (broadcasts) the frame to all other ports on the switch.
2. When the switch receives a unicast frame, it forwards it only to the port to which it is addressed.

A VLAN-aware switch changes the above two rules as follows:

1. When the switch receives a broadcast or multicast frame from a port, it floods the frame **to only those ports that belong to the same VLAN as the frame**.
2. When a switch receives a unicast frame, it forwards it to the port to which it is addressed, **only if the port belongs to the same VLAN as the frame**.
3. A unique number called the VLAN ID identifies each VLAN². It is a 12-bit field in the VLAN tag. You can have a theoretical maximum of 4095 discrete VLANs in your network.³

2. Most switches allow you to assign a name to each VLAN.


3. Some switches support a much smaller number of VLANs. The number of VLANs supported must not be confused with the number of VLAN IDs that can be used. Typically, no limitations exist on which VLAN IDs you can use to identify VLAN groups—most switches support the entire range of the 12-bit value to be used.

Which VLAN Does a Frame Belong To?

The previous section notes that a frame can belong to a VLAN. The next question is—how is this association made?

- A VLAN-aware switch can make the association based on various attributes of the frame (such as Ethernet and IP header content). Example attributes include destination MAC address, IP address, TCP port, Network Layer protocol, and so on.
- Attributes such as “the switch port on which the frame arrived” can also be used. In this case, the switch implicitly assigns a VLAN ID to all frames arriving on a given port.
- A frame can carry explicit VLAN information in a tag that is added to the Ethernet header (explicit VLAN tagging). See Figure 2 for the format of the VLAN tag.

Figure 2 IEEE 802.1Q VLAN Tag in Ethernet Frame



How Does a VLAN-aware Switch Work?

You can configure VLAN-aware switches to add ports to a VLAN group or groups. These switches maintain two simple, related tables:

- a list of ports that belong to each VLAN enabled on the switch
- the set of VLANs enabled on each port

Several varieties of VLAN-aware switches are available:

- The most basic of these switches support port-based VLANs. In a port-based VLAN, the switch port on which the frame arrived determines the VLAN membership of the frame. These switches cannot support more than one VLAN per switch port unless they support VLAN tagging, which is explained in following sections. A simple port-based VLAN that supports VLAN tagging is all you need to implement a VLAN in an HP-UX environment.
- More sophisticated switch offerings enable users to configure VLAN membership rules based on frame content, such as MAC address, TCP/UDP port, IP address, and so on. Doing this can affect switch performance.
- VLAN-aware Layer 3 switches (or routing switches) perform the function of Layer 3 (e.g., IP routing) in addition to VLAN classification.

With regard to other network devices, note the following:

- You can configure an end station to belong to more than one VLAN.
- Shared bandwidth devices, such as hubs, cannot be VLAN aware, though they can be included in a VLAN environment. If a hub is used in a VLAN environment, all nodes on

that hub must belong to the same VLAN or set of VLANs, thereby restricting the benefits of VLANs.

- A common misconception is that because multiple IP subnets can share a single switched infrastructure using VLANs, switching can replace routing in the network. Remember that VLAN is strictly a Data Link Layer (Layer 2) technology. You must use routers for communication between IP subnets, even in a VLAN. [In the case where you are using a VLAN-aware Layer 3 (routing) switch (mentioned previously as one of the available types of VLAN-aware switches), you do not need a separate router, because VLAN-aware routing switches incorporate both the Layer 2 and Layer 3 functions.]

VLAN Tagging

As mentioned previously, you can implement VLAN functionality via explicit frame tagging by switches and end stations. Network switches and end stations that know about VLANs are said to be VLAN aware. Network switches and end stations that can interpret VLAN tags are said to be VLAN tag aware. VLAN-tag-aware switches and end stations add VLAN tags to standard Ethernet frames—a process called explicit tagging. In explicit tagging, the end station or switch determines the VLAN membership of a frame and inserts a VLAN tag in the frame header (see Figure 2), so that downstream link partners can examine just the tag to determine the VLAN membership.

Tagging has several advantages—VLAN association needs to be applied only once at an end station or at an edge switch, so that downstream switches all the way to the destination are relieved of the burden of classifying frames. Tagging at end stations is particularly beneficial because the overhead of frame classification is distributed.

Standards and Interoperability


IEEE 802.1Q specifies the architecture for VLAN tagging—tag format, tag insertion, and tag stripping. The IEEE 802.1Q tag (shown in Figure 2 for an Ethernet frame) also has a provision for priority encoding. The 3-bit priority field in the tagged frame carries priority information. IEEE 802.1p (later incorporated in IEEE 802.1D) has standardized this priority encoding.

VLAN Trunking

Switches that implement only port-based VLAN can support only one VLAN per port. However, if they are tag aware (also called Q-compliant), they can support multiple VLANs per port— one *untagged* VLAN and multiple *tagged* VLANs. If a frame doesn't have an explicit VLAN tag, it is automatically assigned the "untagged VLAN ID" or the "default VLAN ID." An inbound frame that is tagged has its VLAN ID in the frame header. Some switch vendors refer to the ability of handling multiple tagged frames per port as VLAN *trunking*.

If the end node is a VLAN-aware server, the LAN card port on the server can support multiple VLANs per port, as shown in Figure 3.

Figure 3 VLANs Overlapping or Sharing the Same LAN Card Port



VLANs on HP-UX

HP-UX allows users to configure VLAN tagging and association rules at end stations. An efficient implementation of this mechanism has been developed, allowing network administrators to make full use of the advantages of VLANs and VLAN tagging with minimal performance impact.

Features and Advantages

A high-level summary of the features and advantages of using HP-UX VLANs includes the following:

- Host-based IEEE 802.1Q-compliant VLAN tagging
- Supported on HP's PCI and HSC Fast Ethernet and Gigabit Ethernet (1000Base-T and 1000Base-SX) NICs with a free software upgrade (via patches)
- IP subnet-based, protocol-based, and port-based VLAN support
- Support for 802.1p priority encoding
- Configuration using standard HP-UX tools
- IP ToS—802.1p priority conversion
- 1024 VLANs per NIC port
- Designed to work seamlessly with HP Auto Port Aggregation (APA) and HP's high availability products, such as HP Serviceguard
- No changes to network applications are required
- Preserves VLAN configuration across reboot
- Supported on HP-UX 11i

Implementing VLANs on HP-UX

HP-UX implements VLAN tagging via a mechanism called virtual interfaces (VIs). On each NIC port, you may configure multiple VIs, each of which is associated with a unique VLAN ID and 802.1p priority value. Each VI is assigned a virtual PPA (Physical Point of Attachment), which can then be used just like any other PPA—for configuring protocols or attaching to applications, and so on. If you are not familiar with the term PPA, refer to the lan manual page on a system running HP-UX, by running the command `man lan(7)`.

Common Usage Scenarios

Common types of usage scenarios for VLANs on HP-UX include: port-based VLANs, protocol-based VLANs, and IP subnet-based VLANs. Before figuring out which usage scenario

suits your needs, you must understand what each type of usage scenario implies. On HP-UX, the type of VLAN configured on a NIC port depends on how you configure virtual interfaces and use them.

- **Port-based VLAN:** All frames transmitted by a NIC are tagged using only one VLAN ID. The NIC does not transmit or receive any untagged frames.

To implement this on HP-UX, you create just one VI on a given NIC port. All protocols and applications use this virtual interface's virtual PPA to transmit data traffic. Therefore all frames transmitted by that NIC port are tagged with the VLAN ID of that VI.

- **Protocol-based VLAN:** The NIC assigns a unique VLAN ID for each Layer 3 protocol (such as IPv4, IPv6, IPX, and so on). Therefore, the VLAN ID of outbound frames is different for each protocol. An inbound frame is dropped if the protocol and VLAN ID do not match.

To implement this on HP-UX, you create one VI per Layer 3 protocol processed by the NIC. You then configure the protocol (for example, `ifconfig`) using the VPPA of each VI.

- **IP subnet-based VLAN:** The NIC assigns a unique VLAN ID for each IP subnet it belongs to. Therefore, the VLAN ID of outbound frames is different for different destination subnets. An inbound frame is dropped if the IP subnet and VLAN ID do not match.

To implement this on HP-UX, you create one VI per IP subnet. In other words, you first create as many VIs as there are subnets that you want configured on a given NIC port, and then you configure IP addresses on their VPPAs using `ifconfig`.

Determining a Usage Scenario that Meets Your Network Needs

The way you decide to use VLANs in a network depends on the requirements of individual stations in the network. If appropriate, you can even configure all three types of VLANs in a network at the same time. The following are some guidelines for determining which type of VLAN to configure:

- If an end-station NIC needs to belong to only one VLAN, you have two choices:
 - Configure a port-based VLAN on that NIC and enable the corresponding VLAN ID on the switch port to which the NIC is connected. This switch port must be marked "tagged" for that VLAN ID.
 - Keep that end station VLAN unaware. You just need to enable the corresponding VLAN ID on the switch port. This switch port must be marked "untagged".

Typically you need to do this on workstation NICs.

- If an end-station NIC needs to process frames for more than one protocol (such as IPv4, IPv6, or IPX), configure a protocol-based VLAN on that NIC, by assigning one VLAN ID to each protocol. You must also configure the switch port the NIC is connected to with the same VLAN IDs and mark them "tagged" on the switch.
- If an end-station NIC must handle IP packets belonging to multiple subnets, use an IP subnet-based VLAN. Assign a unique VLAN ID to each IP address configured on that NIC. Enable the same VLAN IDs on the switch port to which the NIC is connected, and mark them "tagged" on the switch.
- You can also use combinations of the three types. For example, if your end station processes frames for more than one protocol, and it also serves multiple IP subnets, consider using both protocol-based and subnet-based VLANs.

Priority and Class of Service (CoS)

HP-UX allows you to specify a 3-bit priority encoding (resulting in eight possible values) for each VLAN configured on a NIC port. The VLAN tag carries this value to all the switches on the route. Some switch vendors have implemented a priority mechanism that acts on this 3-bit priority encoded in the VLAN tag (see Figure 2), to provide a rudimentary Class of Service (CoS) differentiated service. For example, in the event of congestion, the switch may give a better

service (for example, lower drop rate or higher scheduling priority) to frames carrying a certain 802.1p priority value in the VLAN tag. For information on priority policies on switches, refer to the switch manufacturer's manuals.

HP-UX allows a user to assign an 802.1p priority to a VLAN. This priority is subsequently encoded in the VLAN tag of the frame's Ethernet header. However, at the time of this writing HP-UX does not enforce any priority mechanisms in either the end station protocol stack, device drivers, or the NICs. In other words, HP-UX end stations do not distinguish between frames with different 802.1p values in the VLAN tag.

IP ToS and 802.1p Conversion—End-to-End Class of Service


HP-UX allows you to map IPv4 Type of Service (ToS) octet to 802.1p priority. The ToS octet is a field in the IP header. Using well-known TCP/IP socket options, applications can specify a desired ToS octet. But since switches are Layer 2 devices, typically they do not look at or act on the priority encoding of the ToS octet. Some switches do, but there may be performance implications. HP-UX VLAN allows IP ToS octet to 802.1p priority conversion. Switches are more likely to implement and enforce 802.1p priority with few or no performance implications because extracting the priority from the VLAN tag is simpler than peeking into the IP header for the required information. Using this mechanism, you can build a network with end-to-end class of service in a LAN.

Typical Customer Configurations

The network shown in Figure 4 depicts a typical usage model. Note the following:

- Sets of workstations are grouped into VLANs, each possibly representing an IP subnet.
- HP-UX servers can be used to serve several VLANs at the same time, with a single point of attachment to the LAN (for example, via a single NIC). This is accomplished by configuring tagged VLANs on the NIC.
- You might want to use IP subnet-based VLANs on your backup server. This is advantageous if you are backing up stations on more than one subnet.
- The servers must be VLAN aware, but the workstations need not be, as they typically tend to belong to a single subnet. If you want to put a workstation NIC on several subnets, you must make it tag aware.
- Using VLAN-aware servers and workstations, you will need only basic VLAN functionality at the switches (for example, port-based VLAN and VLAN-tagging/-trunking capability).

Figure 4 VLAN Implementation Example



A VLAN Implementation Example

This section provides an overview of how to implement an IP subnet-based VLAN in a network with HP-UX. See Figure 4 for reference. The steps include the following:

- Identify the logical partitions in the network. That is, decide how many subnets you want the network to be partitioned into, based on security, performance, and management considerations. Assign VLAN IDs to each subnet.
- Assign subnets to each station in the network. Then configure VLANs on the switches and the end stations as follows:
 - Identify the number of ports needed and implement a monolithic switched LAN infrastructure as shown in Figure 4. The switches must support port-based VLAN and 802.1Q-compliant tagging.⁴
 - Since workstations typically do not belong to more than one subnet, they can be VLAN unaware. You must configure an untagged VLAN on the switch port a VLAN-unaware workstation connects to. That means all frames received from the workstation will be associated with the untagged VLAN. Furthermore, the switch will strip the VLAN tag from all switch-to-end station traffic.
 - Servers typically belong to more than one subnet. Therefore, configure the required number of VLANs on the server NIC, each corresponding to a subnet. In this case, you must configure both the HP-UX server and the switch port it attaches to. On the HP-UX server, create VLANs (virtual interfaces) with appropriate VLAN IDs, and assign IP

4. Typically, in the factory configuration, all ports in the switch are configured with the same default untagged VLAN of VLAN ID 1. This configuration allows a VLAN-aware switch to behave exactly like a traditional LAN switch.

addresses to them. Then configure the same VLANs on the switch port, marking them “tagged.” One (and only one) untagged VLAN can be configured on a switch port.

- If a workstation needs to belong to more than one VLAN and supports tagging, follow the same steps as outlined for the server configuration.

Using HP-UX VLANs with HP Auto Port Aggregation (APA)

HP APA is a software product that creates link aggregates, often called trunks, which provide a logical grouping of two or more physical ports into a single fat pipe. This port arrangement provides more data bandwidth than would otherwise be available. HP APA also allows users to create failover groups, which is a link aggregate in LAN_MONITOR configuration mode, of two or more physical ports. A failover group provides redundancy of network links within the group.

The HP APA interfaces can be in any of the modes, including:

- Manual trunk
- IEEE 802.1d Link Aggregation Control Protocol
- Cisco's Port Aggregation Protocol
- Lan Monitor failover groups



NOTE: The industry leading performance and scalability of the HP APA product are documented in the *HP APA Performance and Scalability White Paper*, located at:

<http://docs.hp.com/en/7662/new-apa-white-paper.pdf>

HP-UX VLAN interfaces are also now configurable over HP APA interfaces. VLAN over APA is supported on HP-UX 11i Version 2 (11.23) with the latest requisite updates and patches, and is supported on HP-UX 11i Version 3 (11.31) as shipped, with no additional patches required.

VLAN over APA server technology offers several benefits. Using VLANs over link aggregates and failover groups:


- Enables network I/O consolidation and higher bandwidth, through efficient usage of multiple links under a single logical APA interface.
- Improves reliability, because the VLANs continue to carry traffic in case the active link failed.
- Provides flexible configuration options for applications, through multiple VLAN interfaces created over highly available APA interfaces. Multiple VLAN interfaces can be configured for separate IP subnets used by different applications. Applications using use-dedicated links prior to consolidation can now use VLAN interfaces created over the highly available HP APA aggregate or LAN Monitor failover group.

Thus, VLAN over APA enables resilient network I/O consolidation on servers with a limited number of slots for network I/O.

Figure 5 illustrates the value proposition of traffic isolation and higher bandwidth realized with VLAN over APA technology. In the figure, each workgroup traffic is isolated from others, but still run over the same set of NICs. The NICs are aggregated for high bandwidth with added resilience to network failures using HP APA LAN Monitor technology⁵.

5. LAN Monitor uses polling packets to monitor the connectivity among the links in the failover group. These packets are always sent untagged. The switch must be configured to allow for the forwarding of these untagged polling packets.

Figure 5 VLAN over APA



Using HP-UX VLANs with HP Virtual Machines (HPVM)

The HPVM product is part of the HP-UX 11i Partitioning Continuum and Virtual Server Environment (VSE) technologies. HPVM is a soft partitioning and virtualization technology that provides operating systems isolation, shared CPU (with sub-CPU granularity), shared I/O, and automatic, dynamic resource allocation that is built in.



NOTE: For additional information about the VSE technologies, see <http://hp.com/go/vse>.

HPVM provides significant flexibility in the sharing of HPVM host network interfaces by one or more HPVM guests virtually connected together through a virtual switch. Each association between an HPVM guest and a virtual switch presents as a virtual network interface on the guest. A virtual switch may be associated with, at most, one host network interface. The host network interface might correspond to a physical NIC or it might be a logical network interface, such as those formed by the HP Auto Port Aggregation (HP APA) product. The HP APA product provides superior performance, scalability and failover capabilities in a network interface by using multiple physical NICs.

The virtual switch, for the most part, emulates a physical Ethernet switch, and allows guests sharing a virtual switch to communicate either with one another, with the HPVM host, or with a remote host on the network. HPVM guests can be connected to one or more virtual switches. Each such connection projects a virtual network I/O device on the guest and represents a link to a virtual port on the virtual switch.

Starting with HPVM version 2.0, the HPVM virtual switch is VLAN aware. The HPVM guests (OS) themselves are still not required to be VLAN aware. The virtual switch virtual ports can be configured for port-based VLANs. Once configured for a VLAN, the HPVM virtual switch will insert the VLAN tag on untagged frames originating from the HPVM guest and strip the VLAN tag on frames destined to the guest. As with a VLAN-aware physical switch, the virtual switch can be used to enforce network traffic isolation based on VLAN memberships of virtual-switch ports.



NOTE: Tagged frames originating from the HPVM guest are not supported and are discarded at the HPVM virtual switch.


Also, only VLAN identifier configuration is allowed over virtual-switch ports. Currently, configuration of 802.1p priorities is not allowed.

There are various scenarios where the system administrator can configure the HPVM virtual switch to utilize the traffic isolation feature of VLAN technology.

In the following example, the HP Integrity Virtual Machines (HPVM) product combines VLAN technology with the high availability, performance and scalability of the HP Auto Port Aggregation technology to enable the HPVM system administrator to enforce isolation policies among guests with resilience and high bandwidth.⁶

In Figure 6 the virtual switch ports 1 and 3 for guests A and C, respectively, are configured for the RED VLAN. Similarly, virtual switch port 2 for guest B is configured for the BLUE VLAN. Port 4 is not configured for any VLAN.

Figure 6 VLANs in an HPVM Configuration



6. When APA is used with HPVM, the load-balancing mode of the link aggregate will always be MAC-address based for HPVM traffic (regardless of the user-configured load-balancing mode).

In the scenario illustrated in Figure 6, the peer-to-peer communication can be between:

- Two HPVM guests virtual network interfaces through the shared virtual switch.
For example, guests A and C (virtual switch ports 1 and 3 respectively) can communicate with each other over the RED VLAN.
- An HPVM guest virtual network interface and the physical network interface on the local HPVM host, through the virtual switch.
For example, guests A and C can communicate to the host VLAN interface, lan5000, over the RED VLAN. Similarly, guest B can communicate with the host VLAN interface lan5001 over the BLUE VLAN.
- An HPVM guest virtual network interface and a remote node through the virtual switch, the associated network interface on the HPVM host and the physical network
For example, guest A can communicate with client X over the RED VLAN. Similarly, guest B can communicate with client Y over the BLUE VLAN.

Also note that the second virtual network interface on guest C (port 4) is connected to a virtual switch port that is not configured for VLANs. Thus, it cannot communicate with either guest A (over the RED VLAN) or with guest B (over the BLUE VLAN). However, it still can communicate with client Z that is also not a member of any VLAN.

The HPVM system administrator can enforce the required isolation policies by configuring appropriate VLAN membership for virtual switch ports. The HPVM system administrator configures virtual switch ports for the same VLAN if communication between corresponding guest interfaces is desired. Note that the VLAN membership rules on the physical network are now extended to HPVM guests by virtue of a VLAN-aware HPVM virtual switch.

Since the VLAN-aware virtual switch extends the VLAN-aware network domain, it follows that VLAN configuration of the HPVM virtual switch(es) must be consistent with the VLAN configuration of the physical network and vice-versa for proper operation and enforcement of policies. For example, if an HPVM guest is to communicate with a remote host over a VLAN, both the HPVM virtual switch port for the guest and the physical switch port that connect the physical interface that backs the virtual switch must be configured for the same VLAN.

Future HP-UX VLAN Feature Additions

HP is investing in the following areas for improvements to the HP-UX VLAN product.

- Generic VLAN Reservation Protocol (GVRP) and Automatic Configuration: GVRP is an IEEE protocol that allows a switch or an end station to advertise its VLAN membership to its link partner. Using this mechanism, we could develop a mechanism for dynamically assigning VLAN membership to end stations, so that you don't need to manually assign VLAN IDs to each NIC on an end station.
- Stack support for 802.1p/Cos/QoS (multi-queues): HP is investigating methods for implementing an end-to-end Class of Service or Quality of Service solution by improving on priority mechanisms in the network stack and NICs. An important component of this solution will be the 802.1p mechanism.
- Application-based VLAN: Application-based VLANs provide the most flexible way for configuring VLANs—VLAN-aware applications determine the membership of the traffic they generate. This mechanism opens up a number of interesting possibilities. For example, a set of stations may negotiate a dynamically created VLAN for the purpose of carrying on a short-term audio or videoconference.
- HP-UX VLAN implementation will be a key value addition to many exciting new technologies in the horizon, such as iSCSI, 10-Gigabit Ethernet, and IPv6.