



an hp white paper

march 2002

planning and  
implementing  
VLANs with hp-ux

## hp-ux VLAN



## table of contents

<b>introduction</b> .....	<b>3</b>
<b>what is VLAN?</b> .....	<b>3</b>
VLAN-aware switches are the key .....	4
<i>which VLAN does a frame belong to?</i> .....	4
<i>how does a VLAN-aware switch work?</i> .....	5
key notes .....	5
typical benefits of VLAN .....	5
VLAN tagging .....	6
standards and interoperability .....	6
VLAN “trunking” .....	6
<b>hp-ux VLANs</b> .....	<b>6</b>
features and advantages at a glance .....	7
types of VLANs supported .....	7
<i>what type of VLAN should I use in my network?</i> .....	7
<i>how do I go about doing this on hp-ux?</i> .....	8
priority and class of service (COS) .....	8
IP ToS and 802.1p conversion—end-to-end class of service .....	9
<b>typical customer configurations</b> .....	<b>9</b>
a VLAN implementation example .....	10
<b>future hp-ux VLAN feature additions</b> .....	<b>11</b>
<b>for more information</b> .....	<b>11</b>


## introduction

The purpose of this white paper is to present network managers with an overview of HP-UX VLAN (Virtual LAN) software for HP servers and workstations. The following HP-UX VLAN, also referred to herein as VLAN, topics will be addressed:

- an introduction to VLAN technology and its benefits—since a VLAN-aware switch is the building block of a VLAN network environment, we first explain how switches implement VLAN
- HP-UX VLANs, features, and advantages—we show how HP-UX fits into a VLAN environment
- planning a VLAN with HP-UX servers and workstations

## what is VLAN?

Virtual LAN (VLAN) technology allows us 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 domain<sup>1</sup> 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.



**figure 1. using VLANs to create multiple independent broadcast domains across switches**

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

**Figure 1** highlights several key differences between traditional LANs and VLAN.

- 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 with VLAN-unaware 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” (i.e., 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.

<sup>1</sup> 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.

## 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.

In order to understand how logical partitioning of a LAN infrastructure is done using VLAN, you should remember the fundamental operation of a traditional switched LAN. Without going into the gory details of switch design, the two simple rules you should 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 VLAN<sup>2</sup>. It is a 12-bit field in the VLAN tag. Therefore you can have a theoretical maximum of 4095 discrete VLANs in your network<sup>3</sup>.

### which VLAN does a frame belong to?

In the previous section we have noted that a frame could belong to a VLAN. The next question is—how is this association made?

1. A VLAN-aware switch can make the association based on various attributes of the frame (e.g., Ethernet and IP header content). Example attributes include destination MAC address, IP address, TCP port, Network Layer protocol, etc.
2. 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.

Finally, 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 exact format of the VLAN tag.

---

<sup>2</sup> Most switches allow you to assign a name to each VLAN.

<sup>3</sup> Some switches only support a much smaller number of VLANs. The number of VLANs supported must not be confused with the number of VLAN IDs that may be used. Typically there are no limitations on which VLAN IDs can be used to identify VLAN groups—most switches allow the entire range of the 12-bit value to be used.

### how does a VLAN-aware switch work?

- VLAN-aware switches can be configured to add ports to a VLAN group or groups. They maintain two simple, related tables: 1) a list of ports that belong to each VLAN enabled on the switch, and 2) the set of VLANs enabled on each port.
- VLAN-aware switches come in many flavors.
  - The most basic VLAN-aware switches support port-based VLANs, meaning that 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 the next sections. As you will see in the next few sections, a simple port-based VLAN that supports VLAN tagging is all that is needed to implement a VLAN in an HP-UX environment.
  - More sophisticated switch offerings allow users to configure VLAN membership rules based on frame content such as MAC address, TCP/UDP port, IP address, etc. But doing this may affect switch performance.
  - VLAN-aware Layer 3 switches (or Routing Switches) will perform the function of Layer 3 (e.g., IP routing) in addition to VLAN classification.

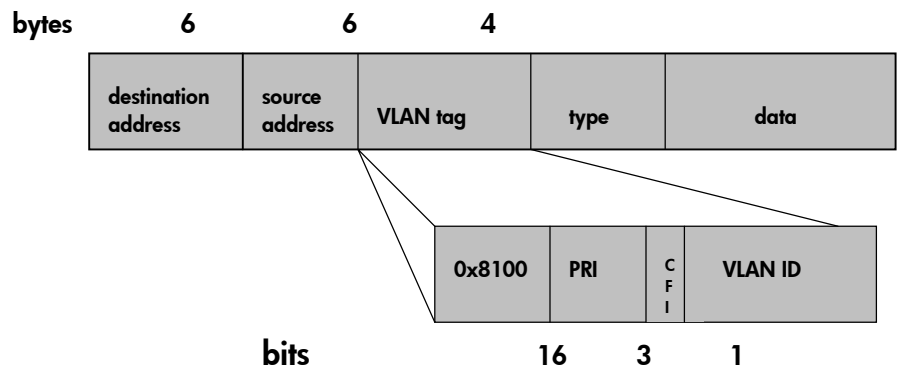
### key notes

- An end station can be configured 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 since 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.

### typical benefits of VLAN

Following are the benefits attributed to VLANs.

- **manageability:** Moves, adds, 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.
- **enhanced security:** Different security domains may be easily constructed to provide various levels of security in the network, since the network design is more flexible than traditional LANs.
- **bandwidth preservation:** A well-designed VLAN will help restrict broadcast and multicast traffic to only those stations that are listening to and responding to the traffic related to that VLAN. The network and computing resources of non-participating stations are unaffected, improving performance.



**figure 2. a tagged Ethernet frame**

## VLAN tagging

As mentioned previously, VLAN functionality may also be implemented via explicit frame tagging by end stations or switches. 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. Devices that can classify frames by inspecting their VLAN tags are called “tag-aware.”

Tagging has several advantages—VLAN association need 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 attractive 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 “PRI” bit in the tagged frame shown in **figure 2** 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 could 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*.

## hp-ux VLANs

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 at a glance

- 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 well known HP-UX tools—*lanadmin* (CLI) and SAM (GUI)
- IP ToS—802.1p priority conversion
- 1024 VLANs per NIC port
- designed to work seamlessly with HP's high availability products, such as MC Serviceguard
- no changes to applications required
- preserve VLAN configuration across reboot
- supported on HP-UX 11i

## types of VLANs supported

HP-UX supports three types of VLANs—port-based, protocol-based, and IP subnet-based.

Before figuring out which type of VLAN suits your needs, you should understand what each type of VLAN means.

1. **port-based VLAN:** All frames transmitted by a NIC are tagged using one and only one VLAN ID. The NIC doesn't transmit or receive any untagged frames.
2. **protocol-based VLAN:** The NIC assigns a unique VLAN ID for each Layer 3 protocol (e.g., IPv4, IPv6, IPX, etc). In other words, the VLAN ID of outbound frames is different for different protocols. An inbound frame is dropped if the protocol and VLAN ID don't match.
3. **IP subnet-based VLAN:** The NIC assigns a unique VLAN ID for each IP subnet it belongs to. In other words, the VLAN ID of outbound frames is different for different destination subnets. An inbound frame is dropped if the IP subnet and VLAN ID don't match.

### what type of VLAN should I use in my network?

The type of VLAN you use depends on the requirements of individual stations in your network. In fact, you can configure all three types of VLANs in your network at the same time. Here are a few suggestions.

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

Typically you will need to do this on workstation NICs.

2. If an HP-UX end station NIC needs to process frames for more than one protocol (e.g., IPv4, IPv6, or IPX), configure a protocol-based VLAN on that NIC—i.e., assign 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.
3. If an end station NIC must handle IP packets belonging to multiple subnets, use IP subnet-based VLAN. Assign a unique VLAN ID to each IP address configured on that NIC. Just as in #2 above, enable the same VLAN IDs on the switch port the NIC is connected to, and mark them “tagged” on the switch.
4. You can also use combinations of #2 and #3 above. For example, if your end station processes frames for more than one protocol, and it also serves multiple IP subnets, consider using both protocol and subnet-based VLANs.

### how do I go about doing this 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, etc. If you are not familiar with the term PPA, please refer to the manual page on a system running HP-UX, by executing the command `man lan(7)`.

The type of VLAN configured on a NIC port depends on how you create and configure virtual interfaces. Here are some examples.

- **port-based VLAN:** 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:** You create one VI per Layer 3 protocol processed by the NIC. You then configure the protocol (e.g., *ifconfig*) using the VPPA of each VI.
- **IP subnet-based VLAN:** 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*.

For more information, please refer to the VLAN user’s guide mentioned in the References section.

### 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 (please 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 (e.g., lower drop rate, higher scheduling priority) to frames carrying a certain 802.1p priority value in the VLAN tag. For more 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.

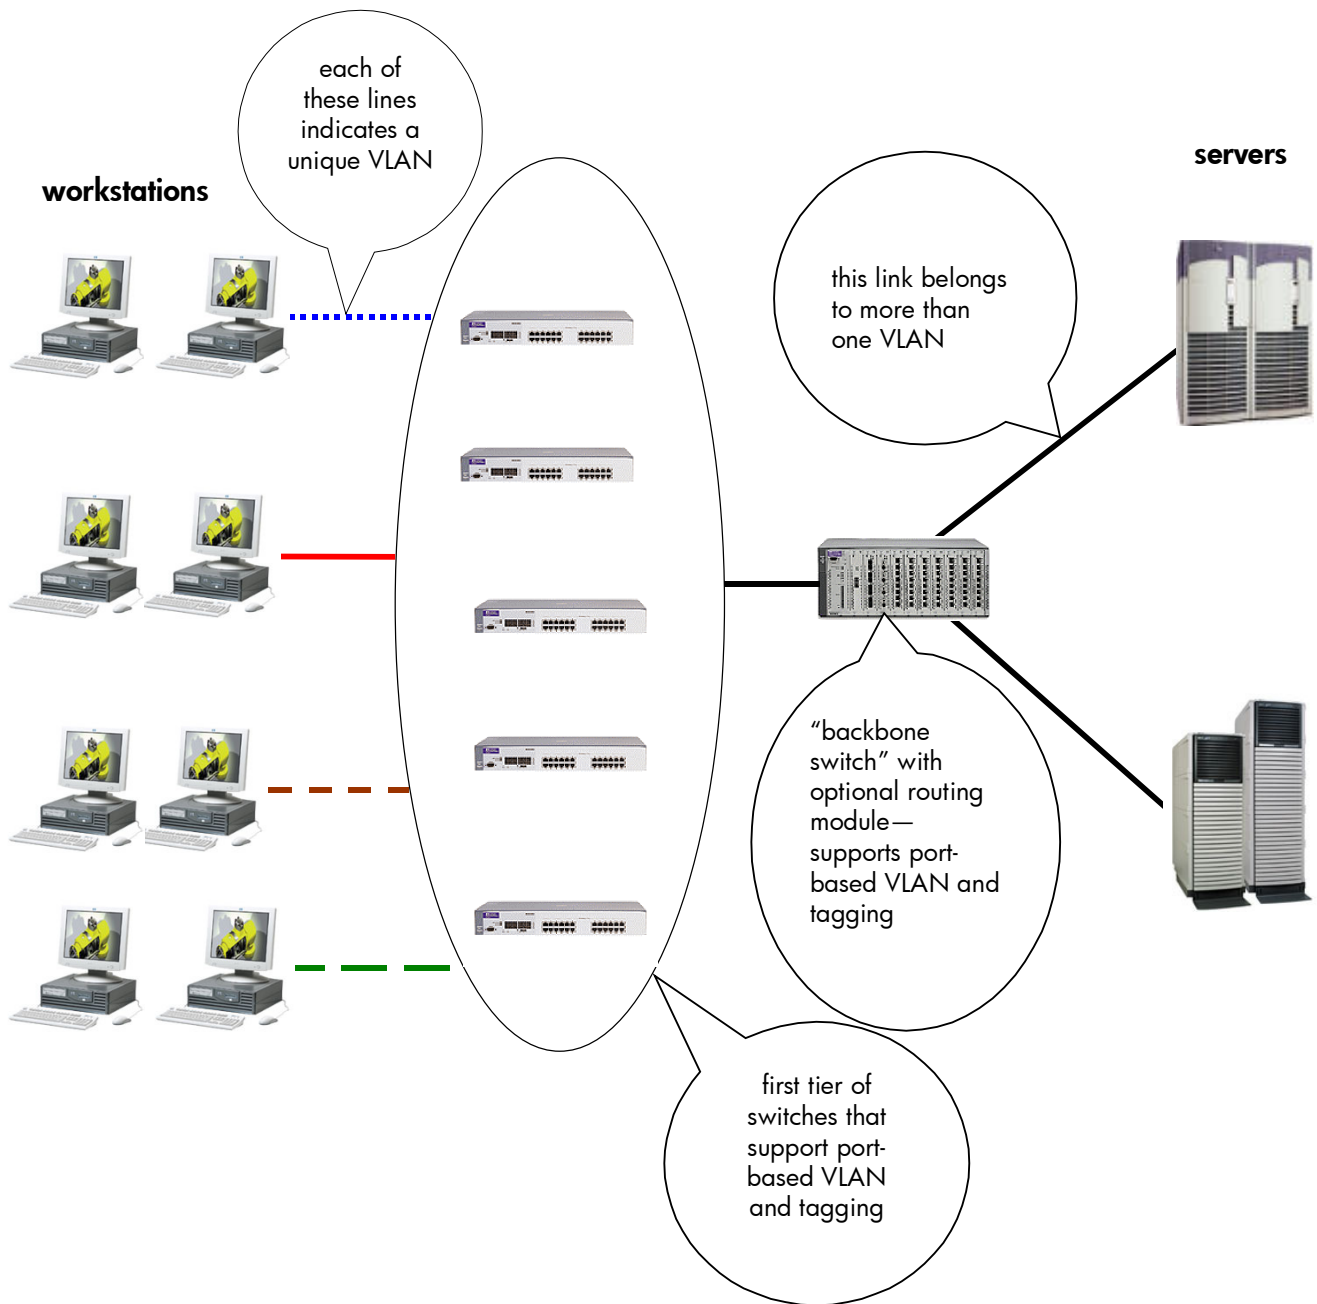
## IP ToS and 802.1p conversion—end-to-end class of service

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/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, we can build a network with end-to-end class of service in a LAN!

## typical customer configurations

The network shown in **figure 3** depicts a general usage model.

- Sets of workstations are grouped into VLANs, each possibly representing an IP subnet.
- HP-UX servers could be used to serve several VLANs at the same time, with a single point of attachment to the LAN (i.e., via a single NIC). This is accomplished by configuring tagged VLANs on the NIC.
- For example, you could 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 desire to put a workstation NIC on several subnets, make it tag-aware.
- Using VLAN-aware servers and/or workstations, you will need only basic VLAN functionality at the switches (i.e., port-based VLAN and VLAN-tagging/-trunking capability).



**figure 3. a VLAN implementation example**

**a VLAN implementation example**

This section gives you an idea as to how you would go about implementing an IP subnet-based VLAN in a network, with HP-UX. Please use **figure 3** for reference.

1. 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.
2. Assign subnets to each station in the network. Then configure VLANs on the switches and the end stations as follows:
  - The switches must support port-based VLAN and 802.1Q-compliant tagging<sup>4</sup>. Identify the number of ports needed and implement a monolithic switched LAN infrastructure as shown in **figure 3**.
  - 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 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.

## future hp-ux VLAN feature additions

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

- Integration of VLAN and auto port aggregation (APA). APA is HP’s Link Aggregation software. Efforts are underway to enable customers to create VLANs on link aggregates.
- 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, IPv6, etc.

## for more information

For additional VLAN information, visit [www.hp.com/go/vlan](http://www.hp.com/go/vlan).

---

<sup>4</sup> Typically, in the factory configuration, all ports in the switch are configured with the same default untagged VLAN (e.g., VLAN ID 1). This configuration allows a VLAN-aware switch to behave exactly like a traditional LAN switch.

Technical information in this document is subject to change without notice.

©Hewlett-Packard Company 2002

Printed in U.S.A.

3/02 rev. 2

5981-0430EN

