Aruba 2530 Advanced Traffic Management Guide for ArubaOS-Switch 16.10

a Hewlett Packard Enterprise company

Part Number: 5200-6744
Published: October 2019
Edition: 1

Notices

The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Links to third-party websites take you outside the Hewlett Packard Enterprise website. Hewlett Packard Enterprise has no control over and is not responsible for information outside the Hewlett Packard Enterprise website.

Acknowledgments

Intel®, Itanium®, Optane®, Pentium®, Xeon®, Intel Inside®, and the Intel Inside logo are trademarks of Intel Corporation in the U.S. and other countries.

Microsoft® and Windows® are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Adobe® and Acrobat® are trademarks of Adobe Systems Incorporated.

Java® and Oracle® are registered trademarks of Oracle and/or its affiliates.

UNIX® is a registered trademark of The Open Group.

This guide provides information on how to configure traffic management features.

Applicable products

This guide applies to these products:

Aruba 2530 Switch Series (J9772A, J9773A, J9774A, J9775A, J9776A, J9777A, J9778A, J9779A, J9780A, J9781A, J9782A, J9783A, J9853A, J9854A, J9855A, J9856A, JL070A)

Switch prompts used in this guide

Examples in this guide are representative and may not match your particular switch/environment. Examples use simplified prompts as follows:

Prompt	Explanation
`switch#`	`#` indicates manager context (authority).
`switch>`	`>` indicates operator context (authority).
`switch(config)#`	`(config)` indicates the config context.
`switch(vlan-x)#`	`(vlan-x)` indicates the vlan context of config, where *x* represents the VLAN ID. For example: `switch(vlan-128)#`.
`switch(eth-x)#`	`(eth-x)` indicates the interface context of config, where `x` represents the interface. For example: `switch(eth-48)#`.
`switch-Stack#`	`Stack` indicates that stacking is enabled.
`switch-Stack(config)#`	`Stack(config)` indicates the config context while stacking is enabled.
`switch-Stack(stacking)#`	`Stack(stacking)` indicates the stacking context of config while stacking is enabled.
`switch-Stack(vlan-x)#`	`Stack(vlan-x)` indicates the vlan context of config while stacking is enabled, where *x* represents the VLAN ID. For example: `switch-Stack(vlan-128)#`.
`switch-Stack(eth-x/y)#`	`Stack(eth-x/y)` indicates the interface context of config, in the form `(eth-<member-in-stack>/<interface>)`. For example: `switch(eth-1/48)#`

Understanding VLANs

Aruba-OS wired switches are 802.1Q VLAN-enabled. In the factory default state, the switch is enabled for up to 256 VLANs. You can reconfigure the switch to support more VLANs. The maximum VLANs allowed varies according to the switch series.

A group of networked ports assigned to a VLAN form a broadcast domain configured on the switch. On a given switch, packets are bridged between source and destination ports that belong to the same VLAN.

VLANs enable grouping users by logical function not physical location. They manage bandwidth usage in networks by:

- Enabling grouping high-bandwidth users on low-traffic segments.
- Organizing users from different LAN segments according to their need for common resources and individual protocols.
- Improving traffic control at the edge of networks by separating traffic of different protocol types.
- Enhancing network security by creating subnets to control in-band access to specific network resources.
- Cross-domain broadcast traffic in the switch is eliminated and bandwidth saved by not allowing packets to flood out all ports.

When configuring VLANs, you will need to plan your VLAN strategy as follows:

Procedure

1. Configure static VLANs with:
   - a name
   - VLAN ID number (VID)
   - port members
2. Include port configuration planning to use dynamic VLANs.
3. Create a map of the logical topology.
4. Create a map of the physical topology.
5. Consider the interaction between VLANs and other features:
   - Spanning Tree Protocol
   - port trunking
   - IGMP
6. Configure at least one VLAN in addition to the default VLAN.
7. Configure all ports that pass traffic for a particular subnet address on the same VLAN.

8. Assign the desired switch ports to the new VLANs.
9. Ensure that the VLAN through which you manage the switch has an IP address, if you are managing VLANs with SNMP in an IP network.

For information on the restrictions when you configure an IP address on a VLAN interface, see the "Comparing port based and protocol based VLAN" table in **Static VLAN operation**.

Static VLAN operation

Static VLANs are configured with a name, VLAN ID number (VID) and port members. For dynamic VLANs, see **GVRP**. 802.1Q compatibility enables you to assign each switch port to multiple VLANs.

Table 1: *Port based and protocol based VLAN*

Function	Port-Based VLANs	Protocol-Based VLANs
IP Addressing	Usually configured with at least one unique IP address. A port-based VLAN can have no IP address. However, this limits switch features available to ports on that VLAN. See "How IP addressing affects switch operation" in the chapter "Configuring IP Addressing" in the *Basic Operation Guide* for the switch. Multiple IP addresses allow multiple subnets within the same VLAN. See the chapter on "Configuring IP Addressing" in the *ArubaOS-Switch Basic Operation Guide* for the switch.	You can configure IP addresses on all protocol VLANs, but IP addressing is used only on IPv4 and IPv6 VLANs. **Restrictions**: Loopback interfaces share the same IP address space with VLAN configurations. The maximum number of IP addresses supported on a switch is 2048; this includes all IP addresses configured for both VLANs and loopback interfaces (except for the default loopback IP address 127.0.0.1). Each IP address configured on a VLAN interface must be unique in the switch; it cannot be used by a VLAN interface or another loopback interface. For more information, see the chapter on "Configuring IP Addressing" in the *ArubaOS-Switch Basic Operation Guide*.
Untagged VLAN Membership	A port can be a member of one untagged, port-based VLAN. All other port-based VLAN assignments for that port must be tagged.	A port can be an untagged member of one protocol VLAN of a specific protocol type, such as IPX or IPv6. If the same protocol type is configured in multiple protocol VLANs, then a port can be an untagged member of only one of those. For example, if you have two protocol VLANs, 100 and 200 and both include IPX, then a port can be an untagged member of either VLAN 100 or VLAN 200, but not both. A port's untagged VLAN memberships can include up to four different protocol types. It can be an untagged member of one of the following: • Four single-protocol VLANs • Two protocol VLANs where one VLAN includes a single protocol and the other includes up to three protocols • One protocol VLAN where the VLAN includes four protocols.

Table Continued

Function	Port-Based VLANs	Protocol-Based VLANs
Tagged VLAN Membership	A port can be a tagged member of any port-based VLAN.	A port can be a tagged member of any protocol-based VLAN.
Routing	If the switch configuration enables IP routing, the switch can internally route IP (IPv4) traffic between port-based VLANs and between port-based and IPv4 protocol-based VLANs. If the switch is not configured to route traffic internally between port-based VLANs, then an external router must be used to move traffic between VLANs.	If the switch configuration enables IP routing, the switch can internally route IPv4 traffic as follows: • Between multiple IPv4 protocol-based VLANs • Between IPv4 protocol-based VLANs and port-based VLANs. Other protocol-based VLANs require an external router for moving traffic between VLANs. **NOTE:** NETbeui and SNA are non-routable protocols. End stations intended to receive traffic in these protocols must be attached to the same physical network.

Commands for Configuring Static VLANs	`vlan <vid> {tagged | untagged <port-list>}`	`vlan <vid> protocol {ipx | ipv4 | ipv6 | arp | appletalk | sna | netbeui}`
		`vlan <vid> {tagged | untagged <port-list>}`

VLAN environments

You can configure different VLAN types in any combination. The default VLAN will always be present. For more on the default VLAN, see **VLAN support and the default VLAN**.

VLAN environment	Elements
The default VLAN (port-based; VID of 1) only	In the default VLAN configuration, all ports belong to VLAN 1 as untagged members. VLAN 1 is a port-based VLAN.
Multiple VLAN environment	In addition to the default VLAN, the configuration can include one or more other port-based VLANs and one or more protocol VLANs. The maximum VLANs allowed on a switch vary according to the switch. For details on the maximum VLANs allowed for your switch, see **Changing the number of VLANs allowed on the switch (CLI)** on page 44. UsingVLAN tagging, ports can belong to multiple VLANs of all types.Enabling routing on the switch enables it to route IPv4 and IPv6 traffic between port-based VLANs and between port-based VLANs and IPv4 protocol VLANs. Routing other types of traffic between VLANs requires an external router capable of processing the appropriate protocols.

VLAN operation

General VLAN operation

- A VLAN is composed of multiple ports operating as members of the same subnet or broadcast domain.
- Ports on multiple devices can belong to the same VLAN.
- Traffic moving between ports in the same VLAN is bridged (or switched).
- Traffic moving between different VLANs must be routed.
- A static VLAN is an 802.1Q-compliant VLAN, configured with one or more ports that remain members regardless of traffic usage.
- A dynamic VLAN is an 802.1Q-compliant VLAN membership that the switch temporarily creates on a port to provide a link to another port either in the same VLAN on another device.

Types of static VLANs available in the switch

Port-based VLANs

This type of static VLAN creates a specific layer-2 broadcast domain comprised of member ports that bridge traffic among themselves. Port-Based VLAN traffic is routable on the switches covered in this guide.

Protocol-based VLANs

This type of static VLAN creates a layer-3 broadcast domain for traffic of a particular protocol and is composed of member ports that bridge traffic of the specified protocol type among themselves. Some protocol types are routable on the switches covered in this guide.

Designated VLANs

The switch uses these static, port-based VLAN types to separate switch management traffic from other network traffic. While these VLANs are not limited to management traffic, they provide improved security and availability.

Default VLAN:

This port-based VLAN is always present in the switch and, in the default configuration, includes all ports as members. See **VLAN support and the default VLAN** on page 53.

Except for an IP address and subnet, no configuration steps are needed.

A switch in the default VLAN configuration

In this example, devices connected to these ports are in the same broadcast domain.

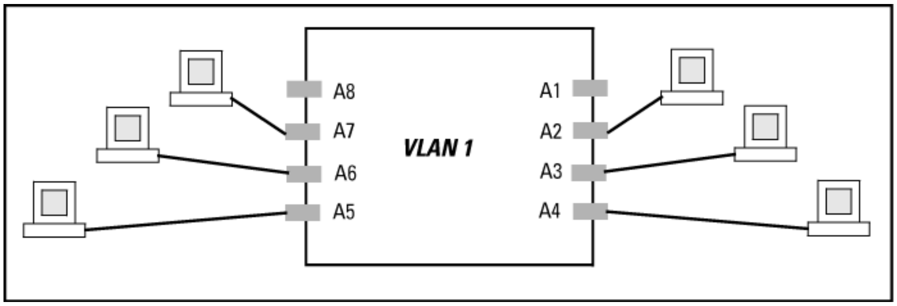

Primary VLAN:

The switch uses this port-based VLAN to run certain features and management functions, including DHCP/ Bootp responses for switch management. In the default configuration, the Default VLAN is also the Primary VLAN. However, any port-based, non-default VLAN can be designated the Primary VLAN. See **The primary VLAN** on page 53.

Secure Management VLAN:

This optional, port-based VLAN establishes an isolated network for managing switches that support this feature. Access to this VLAN and to the switch's management functions are available only through ports configured as members. See **The primary VLAN** on page 53.

Voice VLANs:

This optional, port-based VLAN type enables separating, prioritizing, and authenticating voice traffic moving through your network, avoiding the possibility of broadcast storms affecting VoIP Voice-over-IP) operation. See **Using voice VLANs** on page 52.

NOTE: In a multiple-VLAN environment that includes older switch models there may be problems related to the same MAC address appearing on different ports and VLANs on the same switch. In such cases, the solution is to impose cabling and VLAN restrictions. For more on this topic, see **Multiple VLAN considerations** on page 27.

Multiple port-based VLANs

In the following example, routing within the switch is disabled (the default). Thus, communication between any routable VLANs on the switch must go through the external router. In this case, VLANs W and X can exchange traffic through the external router, but traffic in VLANs Y and Z is restricted to the respective VLANs.

VLAN 1(the default) is present but not shown. The default VLAN cannot be deleted from the switch, but ports assigned to other VLANs can be removed from the default VLAN. If internal (IP) routing is enabled on the switch, then the external router is not needed for traffic to move between port-based VLANs.

A switch with multiple VLANs configured and internal routing disabled

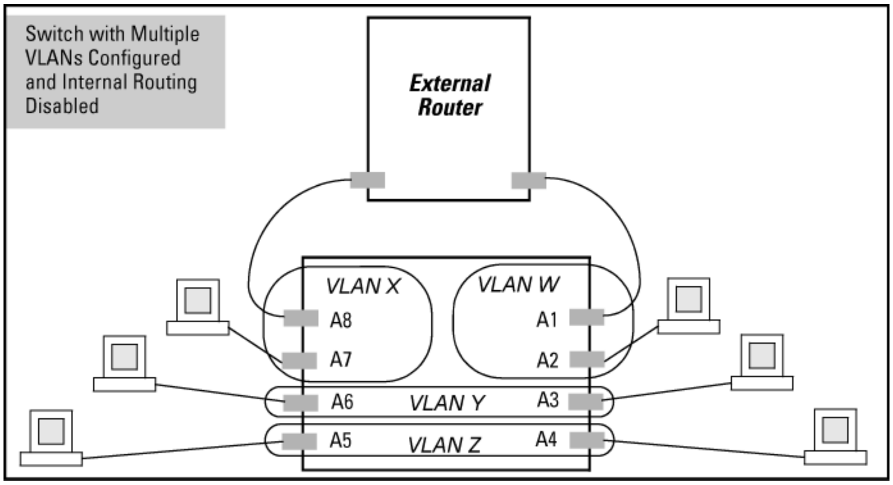
Switch with Multiple VLANs Configured and Internal Routing Disabled
External Router
VLAN X
VLAN W
A8
A7
A1
A2
A6
VLAN Y
A3
A5
VLAN Z
A4

Protocol VLAN environment

The figure in **Multiple port-based VLANs** illustrates a protocol VLAN environment also. In this case, VLANs W and X represent routable protocol VLANs. VLANs Y and Z can be any protocol VLAN.

As noted for the discussion of multiple port-based VLANs, VLAN 1 is not shown. Enabling internal (IP) routing on the switch allows IP traffic to move between VLANs on the switch, but routable, non-IP traffic always requires an external router.

Routing options for VLANs

Table 2: *Options for routing between VLAN types in the switch*

Note that SNA and NETbeui are not routable protocol types. End stations intended to receive traffic in these protocols must be attached to the same physical network.

		Port-Based	IPX	IPv4	IPv6	ARP	AppleTalk	SNA	NETbeui
Port-Based		Yes	—	Yes	—	—	—	—	—
Protocol	IPX	—	Yes	—	—	—	—	—	—
	IPX4	Yes	—	Yes	—	—	—	—	—
	IPV6	—	—	—	Yes[1]	—	—	—	—
	ARP	—	—	—	—	Yes[1]	—	—	—
	AppleTalk	—	—	—	—	—	Yes[1]	—	—
	SNA	—	—	—	—	—	—	—	—
	NETbeui	—	—	—	—	—	—	—	—

802.1Q VLAN tagging

A port can be a member of more than one VLAN of the same type if the device to which the port connects complies with the 802.1Q VLAN standard.

For example, a port connected to a central server using a network interface card (NIC) that complies with the 802.1Q standard can be a member of multiple VLANs, allowing members of multiple VLANs to use the server.

- Although these VLANs cannot communicate with each other through the server, they can all access the server over the same connection from the switch.
- Where VLANs overlap in this way, VLAN "tags" are used in the individual packets to distinguish between traffic from different VLANs.
- A VLAN tag includes the particular VLAN ID. (VID) of the VLAN on which the packet was generated.

For more on this topic, see **Configuring or changing static VLAN per-port settings (CLI)** on page 34.

Overlapping VLANs using the same server

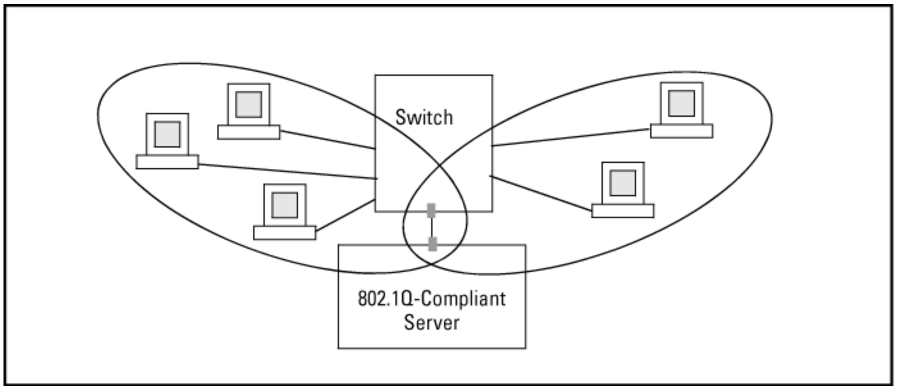

Similarly, using 802.1Q-compliant switches, you can connect multiple VLANs through a single switch-to-switch link.

Connecting multiple VLANs through the same link

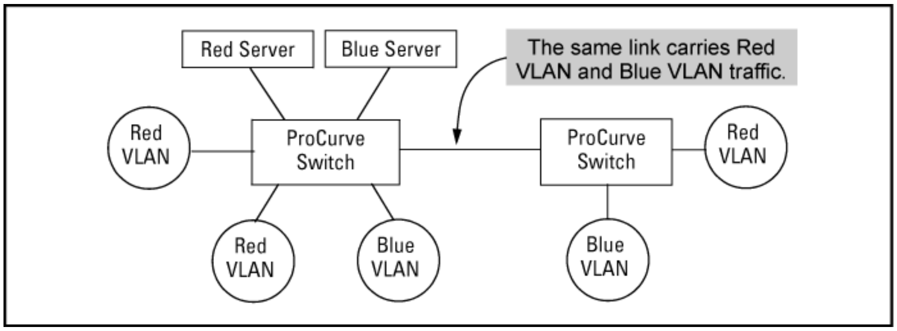

Introducing tagged VLANs into legacy networks running only untagged VLANs

You can introduce 802.1Q-compliant devices into networks that have built untagged VLANs based on earlier VLAN technology. The fundamental rule is that legacy/untagged VLANs require a separate link for each VLAN, while 802.1Q, or tagged VLANs can combine several VLANs in one link. Thus on the 802.1Q-compliant device, separate ports (configured as untagged) must be used to connect separate VLANs to non-802.1Q devices.

Tagged and untagged VLAN technology in the same network

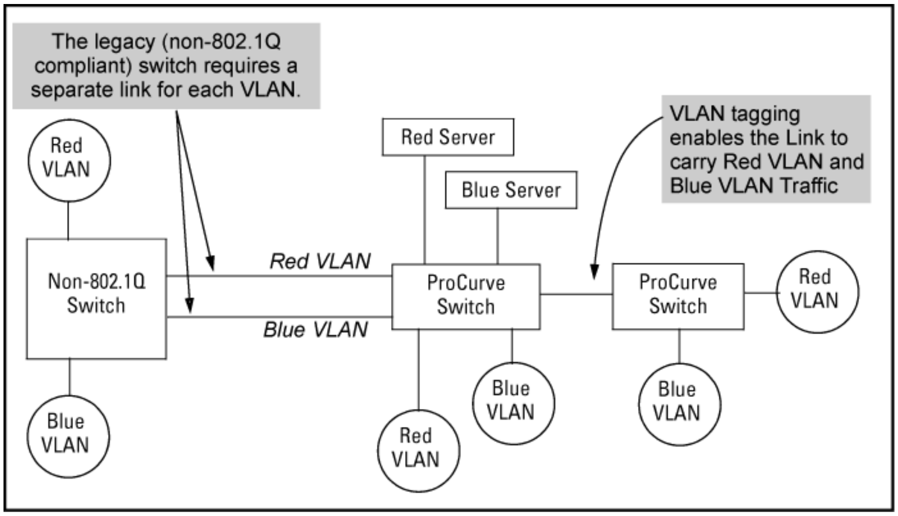
The legacy (non-802.1Q compliant) switch requires a separate link for each VLAN.
Red VLAN
Red Server
Blue Server
VLAN tagging enables the Link to carry Red VLAN and Blue VLAN Traffic
Non-802.1Q Switch
Red VLAN
ProCurve Switch
ProCurve Switch
Red VLAN
Blue VLAN
Blue VLAN
Blue VLAN
Blue VLAN
Red VLAN

VLAN tagging rules

When tagging is needed

When a port belongs to two or more VLANs of the same type, they remain as separate broadcast domains and cannot receive traffic from each other without routing.

NOTE:

If multiple, non-routable VLANs exist in the switch—such as NETbeui protocol VLANs—they cannot receive traffic from each other.

Inbound tagged packets

The switch requires VLAN tagging on a given port if the port will be receiving inbound, tagged VLAN traffic that should be forwarded.

If a tagged packet arrives on a port that is not a tagged member of the VLAN indicated by the packet's VID, the switch drops the packet.

Similarly, the switch drops an inbound, tagged packet if the receiving port is an untagged member of the VLAN indicated by the packet's VID.

Untagged packet forwarding

If the only authorized, inbound VLAN traffic on a port arrives untagged, then the port must be an untagged member of that VLAN. This is the case where the port is connected to a non-802.1Q compliant device or is assigned to only one VLAN.

To enable an inbound port to forward an untagged packet, the port must be an untagged member of either a protocol VLAN matching the packet's protocol, or an untagged member of a port-based VLAN.

That is, when a port receives an incoming, untagged packet, it processes the packet according to the following ordered criteria:

1. If the port has no untagged VLAN memberships, the switch drops the packet.
2. If the port has an untagged VLAN membership in a protocol VLAN that matches the protocol type of the incoming packet, then the switch forwards the packet on that VLAN.
3. If the port is a member of an untagged, port-based VLAN, the switch forwards the packet to that VLAN. Otherwise, the switch drops the packet.

Figure 1: *Untagged VLAN operation*

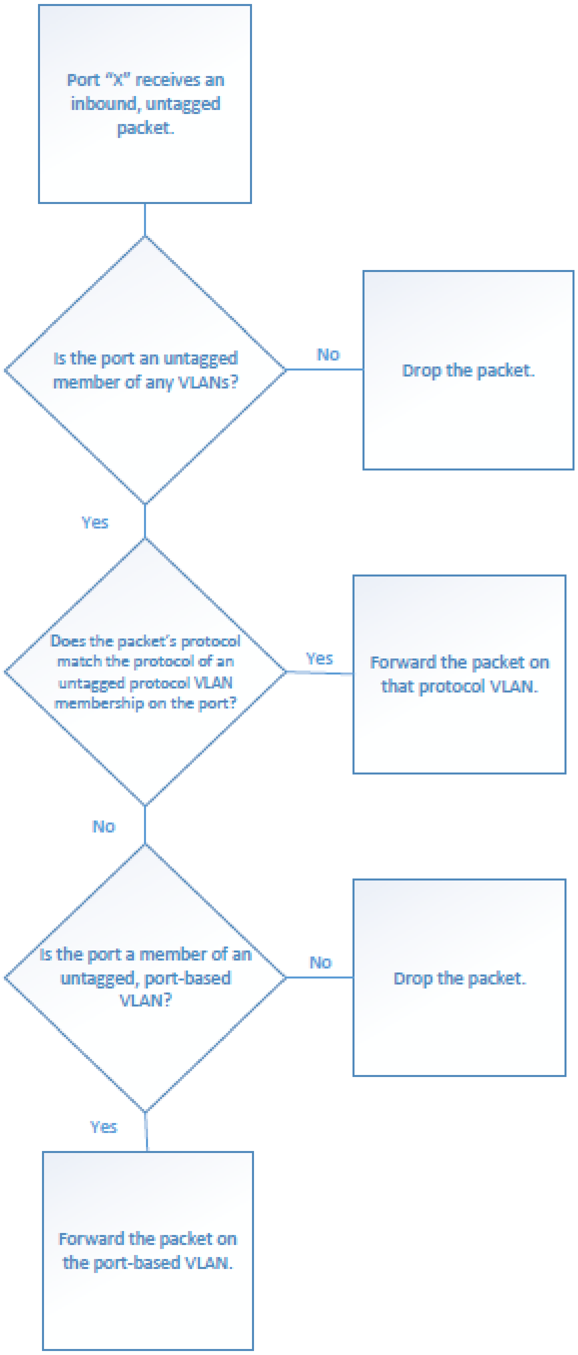
Port "X" receives an inbound, untagged packet.
Is the port an untagged member of any VLANs?
No
Drop the packet.
Yes
Does the packet's protocol match the protocol of an untagged protocol VLAN membership on the port?
Yes
Forward the packet on that protocol VLAN.
No
Is the port a member of an untagged, port-based VLAN?
No
Drop the packet.
Yes
Forward the packet on the port-based VLAN.

Tagged packet forwarding

If a port is a tagged member of the same VLAN as an inbound, tagged packet received on that port, then the switch forwards the packet to an outbound port on that VLAN.

To enable the forwarding of tagged packets, any VLAN to which the port belongs as a tagged member must have the same VID as that carried by the inbound, tagged packets generated on that VLAN.

Figure 2: *Tagged VLAN operation*

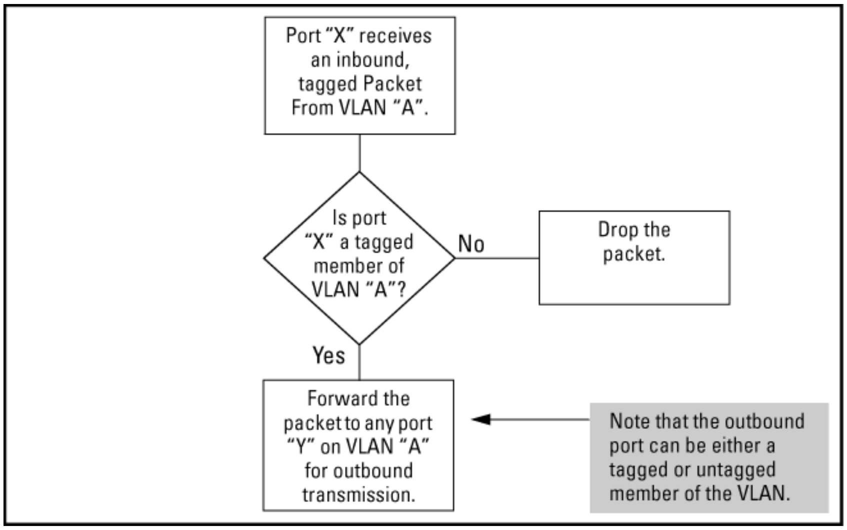
Port "X" receives an inbound, tagged Packet From VLAN "A".
Is port "X" a tagged member of VLAN "A"?
No
Drop the packet.
Yes
Forward the packet to any port "Y" on VLAN "A" for outbound transmission.
Note that the outbound port can be either a tagged or untagged member of the VLAN.

See also **Multiple VLAN considerations** on page 27.

CAUTION: Rate limiting may behave unpredictably on a VLAN if the VLAN spans multiple modules or port-banks.

This also applies if a port on a different module or port-bank is added to an existing VLAN. Hewlett Packard Enterprise does not recommend configuring rate limiting on VLANs that include ports spanning modules or port-banks.

In the following example, ports 2, 3 and 24 form one VLAN, with ports 1 through 24 in the same port-bank. Ports 28, 29 and 32 form a second VLAN. These ports are also in the same port-bank, which includes ports 25 through 48. Rate limiting will operate as expected for these VLANs.

Figure 3: *Example of VLANs using ports from the same port-bank for each VLAN*

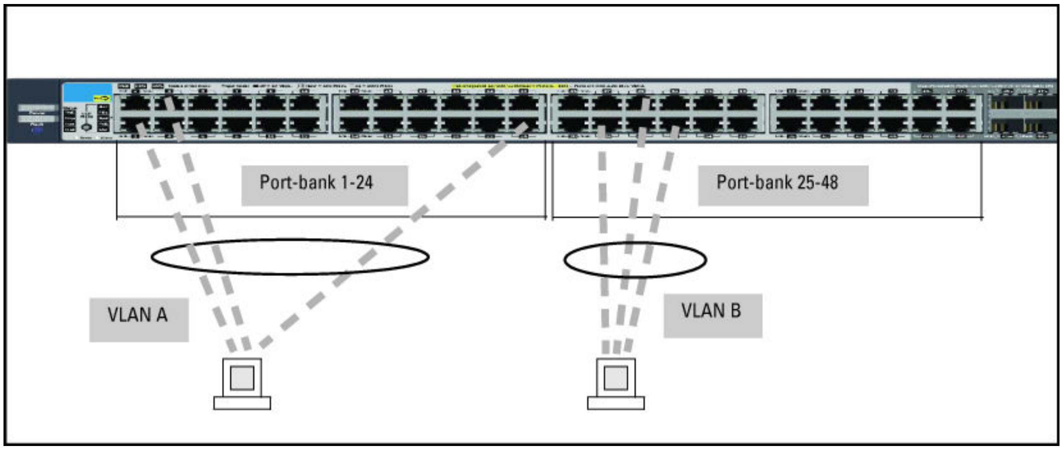

Applying VLAN tagging

Example of tagged and untagged VLAN port assignments

If port 7 on an 802.1Q-compliant switch is assigned to only the Red VLAN, the assignment can remain "untagged" because the port will forward traffic only for the Red VLAN. However, if both the Red and Green VLANs are assigned to port 7, then at least one of those VLAN assignments must be "tagged" so that Red VLAN traffic can be distinguished from Green VLAN traffic.

Figure 4: *Tagged and untagged VLAN port assignments*

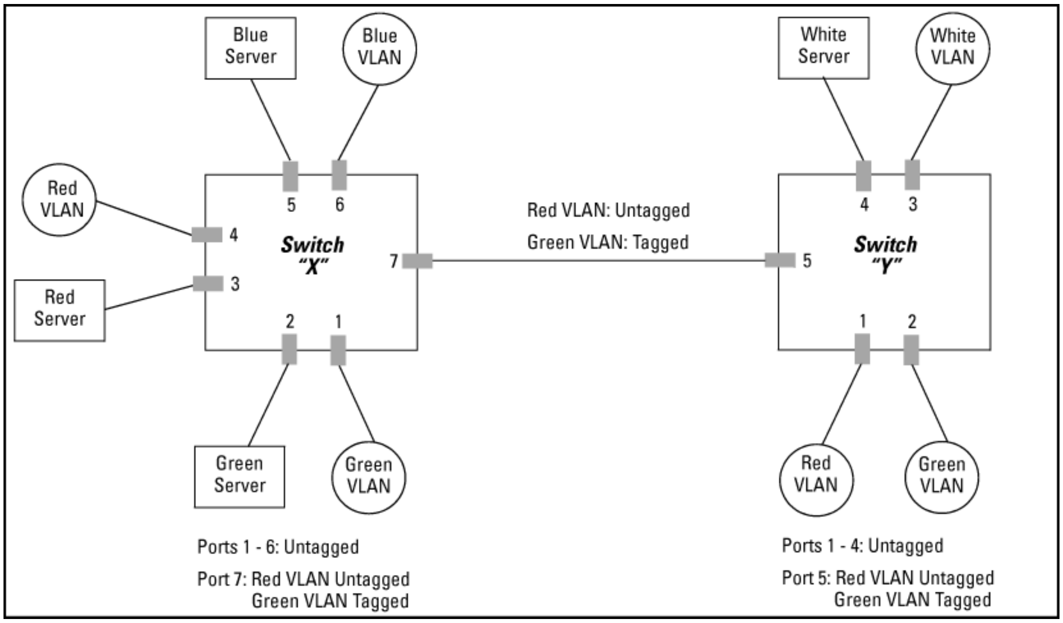

In switch X:

- VLANs assigned to ports X1 - X6 can be untagged because there is only one VLAN assignment per port. Red VLAN traffic will go out only the Red ports, Green VLAN traffic will go out only the Green ports, and so on. Devices connected to these ports do not have to be 802.1Q-compliant.
- However, because both the Red VLAN and the Green VLAN are assigned to port X7, at least one of the VLANs must be tagged for this port.

In switch Y:

- VLANs assigned to ports Y1 - Y4 can be untagged because there is only one VLAN assignment per port. Devices connected to these ports do not have to be 802.1Q-compliant.
- Because both the Red VLAN and the Green VLAN are assigned to port Y5, at least one of the VLANs must be tagged for this port.

In both switches:

The ports on the link between the two switches must be configured the same. As shown in the following figure, the Red VLAN must be untagged on port X7 and Y5 and the Green VLAN must be tagged on port X7 and Y5, or the opposite way.

NOTE: Each 802.1Q-compliant VLAN must have its own unique VID number and that VLAN must be given the same VID in every device where configured. That is, if the Red VLAN has a VID of 10 in switch X, then 10 must also be the Red VID in switch Y.

Figure 5: *Example of VLAN ID numbers assigned in the VLAN names screen*

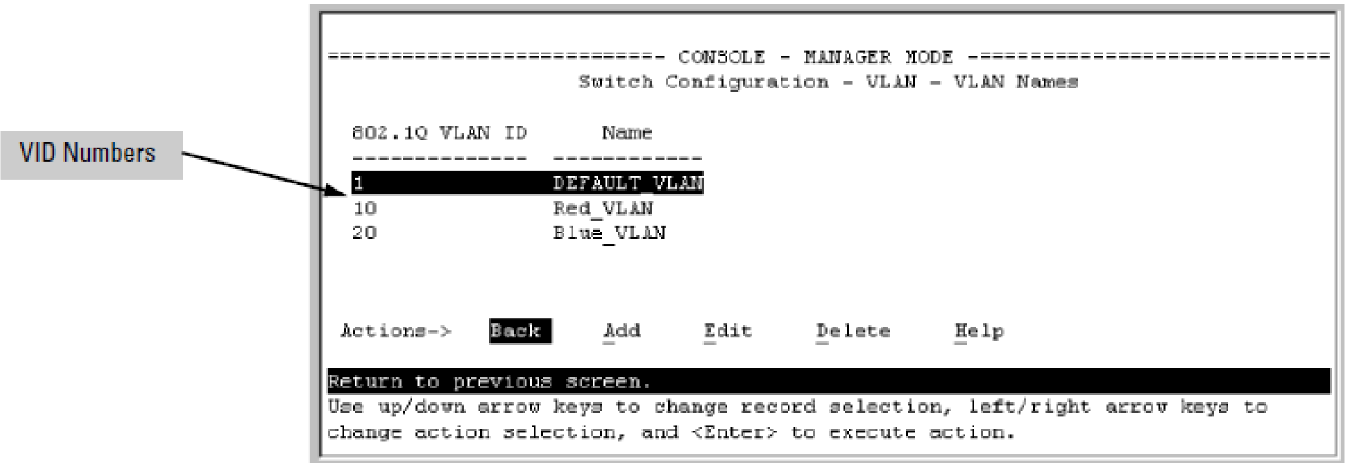

```
==========================- CONSOLE - MANAGER MODE -============================
                          Switch Configuration - VLAN - VLAN Names

  802.1Q VLAN ID      Name
  --------------   ------------
  1                DEFAULT_VLAN
  10               Red_VLAN
  20               Blue_VLAN


 Actions->   Back     Add      Edit     Delete     Help

Return to previous screen.
Use up/down arrow keys to change record selection, left/right arrow keys to
change action selection, and <Enter> to execute action.
```

Additional VLAN tagging considerations

- Since the purpose of VLAN tagging is to allow multiple VLANs on the same port, any port that has only one VLAN assigned to it can be configured as "Untagged" (the default) if the authorized inbound traffic for that port arrives untagged.
- Any port with two or more VLANs of the same type can have one such VLAN assigned as "Untagged." All other VLANs of the same type must be configured as "Tagged," that is:

Port-Based VLANs	Protocol VLANs
A port can be a member of one untagged, port-based VLAN. All other port-based VLAN assignments for that port must be tagged.	A port can be an untagged member of one protocol-based VLAN of each protocol type. When assigning a port to multiple, protocol-based VLANs sharing the same type, the port can be an untagged member of only one such VLAN.
A port can be a tagged member of any port-based VLAN.	A port can be a tagged member of any protocol-based VLAN. See above.
A given VLAN must have the same VID on all 802.1Q-compliant devices in which the VLAN occurs. Also, the ports connecting two 802.1Q devices should have identical VLAN configurations.	

- If all end nodes on a port comply with the 802.1Q standard and are configured to use the correct VID, you can configure all VLAN assignments on a port as "Tagged" if doing so either makes it easier to manage your VLAN assignments, or if the authorized, inbound traffic for all VLANs on the port will be tagged.

For a summary and flowcharts of untagged and tagged VLAN operation on inbound traffic, see the following under **VLAN tagging rules** on page 21:

- **"Inbound Tagged Packets"**
- "Untagged Packet Forwarding" and **Figure 1: Untagged VLAN operation** on page 22
- "Tagged Packet Forwarding" and **Figure 2: Tagged VLAN operation** on page 23

Example of Networked 802.1Q-compliant devices with multiple VLANs on some ports

In the following network, switches X and Y and servers S1, S2, and the AppleTalk server are 802.1Q-compliant. (Server S3 could also be 802.1Q-compliant, but it makes no difference for this example.) This network includes both protocol-based (AppleTalk) VLANs and port-based VLANs.

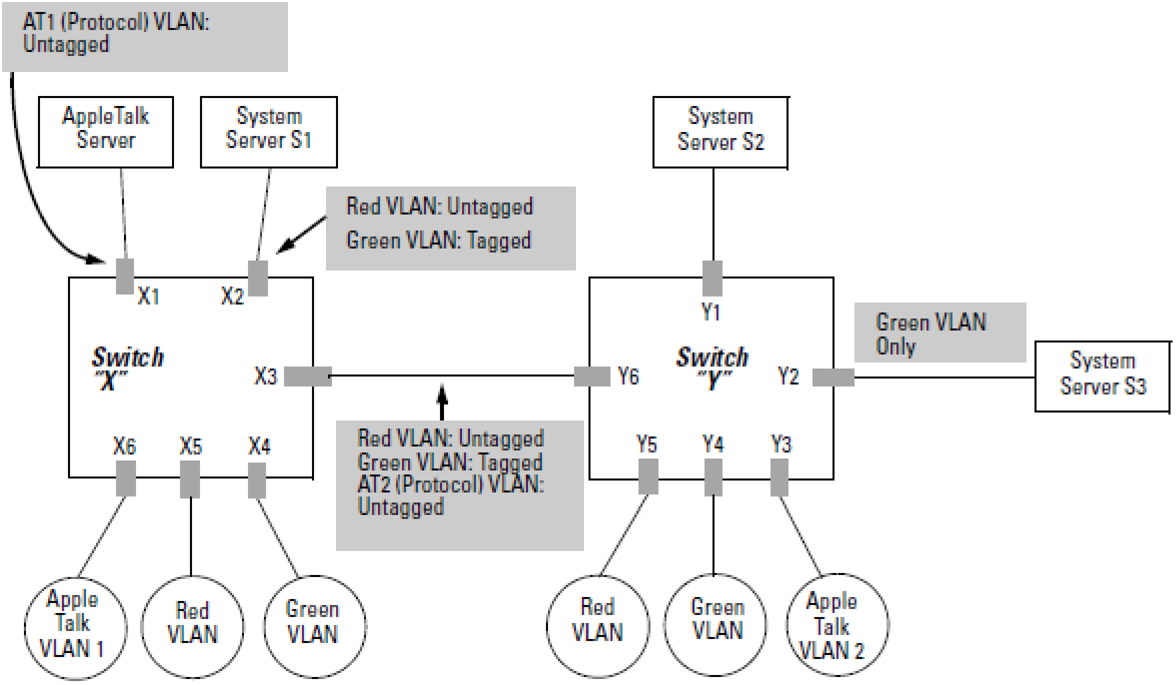

- The VLANs assigned to ports X4 - X6 and Y2 - Y5 can all be untagged because there is only one VLAN assigned per port.
- Port X1 has two AppleTalk VLANs assigned, which means that one VLAN assigned to this port can be untagged and the other must be tagged.
- Ports X2 and Y1 have two port-based VLANs assigned, so one can be untagged and the other must be tagged on both ports.
- Ports X3 and Y6 have two port-based VLANs and one protocol-based VLAN assigned. Thus, one port-based VLAN assigned to this port can be untagged and the other must be tagged. Also, since these two ports share the same link, their VLAN configurations must match.

In the table, "No" means that the port is not a member of that VLAN. For example, port X3 is not a member of the Red VLAN and does not carry Red VLAN traffic. Also, if GVRP were enabled (port-based only), Auto would appear instead of No.

Switch X					Switch Y				
Port	AT-1 VLAN	AT-2 VLAN	Red VLAN	Green VLAN	Port	AT-1 VLAN	AT-2 VLAN	Red VLAN	Green VLAN
X1	Untagged	Tagged	No	No	Y1	No	No	Untagged	Tagged
X2	No	No	Untagged	Tagged	Y2	No	No	No	Untagged

Table Continued

Switch X					Switch Y				
Port	**AT-1 VLAN**	**AT-2 VLAN**	**Red VLAN**	**Green VLAN**	**Port**	**AT-1 VLAN**	**AT-2 VLAN**	**Red VLAN**	**Green VLAN**
X3	No	Untagged	Untagged	Tagged	Y3	No	Untagged	No	No
X4	No	No	No	Untagged	Y4	No	No	No	Untagged
X5	No	No	Untagged	No	Y5	No	No	Untagged	No
X6	Untagged	No	No	No	Y6	No	Untagged	Untagged	Tagged

NOTE: VLAN configurations on ports connected by the same link must match. Because ports X2 and Y5 are opposite ends of the same point-to-point connection, both ports must have the same VLAN configuration, configuring the Red VLAN as "Untagged" and the Green VLAN as "Tagged."

Multiple VLAN considerations

Switches use a forwarding database to maintain awareness of which external devices are located on which VLANs. Some switches, such as the switches covered in this guide, have a multiple forwarding database, which means the switch allows multiple database entries of the same MAC address, with each entry showing the (different) source VLAN and source port. Other switch models have a single forwarding database, which allows only one database entry of a unique MAC address, along with the source VLAN and source port on which it is found. All VLANs on a switch use the same MAC address. Thus, connecting a multiple forwarding database switch to a single forwarding database switch where multiple VLANs exist imposes some cabling and port VLAN assignment restrictions. The following table illustrates the functional difference between the two database types.

Table 3: *Forwarding database content*

Multiple forwarding database			Single forwarding database		
MAC address	**Destination VLAN ID**	**Destination port**	**MAC address**	**Destination VLAN ID**	**Destination port**
0004ea-84d9f4	1	A5	0004ea-84d9f4	100	A9
0004ea-84d9f4	22	A12	0060b0-880af9	105	A10
0004ea-84d9f4	44	A20	0060b0-880a81	107	A17
0060b0-880a81	33	A20			
This database allows multiple destinations for the same MAC address. If the switch detects a new destination for an existing MAC entry, it just adds a new instance of that MAC to the table.			This database allows only one destination for a MAC address. If the switch detects a new destination for an existing MAC entry, it replaces the existing MAC instance with a new instance showing the new destination.		

All switches covered in this guide use a multiple forwarding database.

Single forwarding database operation

When a packet arrives with a destination MAC address that matches a MAC address in the switch's forwarding table, the switch tries to send the packet to the port listed for that MAC address. But if the destination port is in a different VLAN than the VLAN on which the packet was received, the switch drops the packet. This is not a problem for a switch with a multiple forwarding database because the switch allows multiple instances of a given MAC address, one for each valid destination. However, a switch with a single forwarding database allows only one instance of a given MAC address.

TIP: If you connect both switch types through multiple ports or trunks belonging to different VLANs and enable routing on the switch with the multiple-forwarding database, then the port and VLAN record maintained on the switch with the single-forwarding database for the multiple-forwarding database can change frequently. This may cause poor performance and the appearance of an intermittent or broken connection.

Switch performance is unreliable

The following example provides a method to identify and correct an unsupported configuration.

Symptom

Poor switch performance, unreliable switch performance, dropped packets, discarded packets, appearance of intermittent or broken links.

Cause

Incorrect switch configuration.

As shown in the following figure, two switches are connected using two ports on each, and the MAC address table for Switch A will sometimes record the switch as accessed on port A1 (VLAN 1) and at other times as accessed on port B1 (VLAN 2).

Procedure

1. **PC A** sends an IP packet to **PC B**.
2. The packet enters VLAN 1 in the switch with the MAC address of the switch in the destination field. Because the switch has not yet learned this MAC address, it does not find the address in its address table and floods the packet out all ports, including the VLAN 1 link (port A1) to the switch. The switch then routes the packet through the VLAN 2 link to the switch, which forwards the packet on to PC B. Because the switch received the packet from the switch on VLAN 2 (port B1), the switch's single forwarding database records the switch as being on port B1 (VLAN 2).
3. **PC A** now sends a second packet to **PC B**. The packet again enters VLAN 1 in the switch with the MAC address of the switch in the destination field. However, this time the switch's single forwarding database indicates that the switch is on port B1 (VLAN 2) and the switch **drops** the packet instead of forwarding it.
4. Later, the switch transmits a packet to the switch through the VLAN 1 link and the switch updates its address table to show that the switch is on port A1 (VLAN 1) instead of port B1 (VLAN 2). Thus, the switch's information on the location of the switch **changes over time**, and the switch discards some

packets directed through it for the switch. This causes poor performance and the appearance of an intermittent or broken link.

Figure 6: *Invalid forwarding configuration*

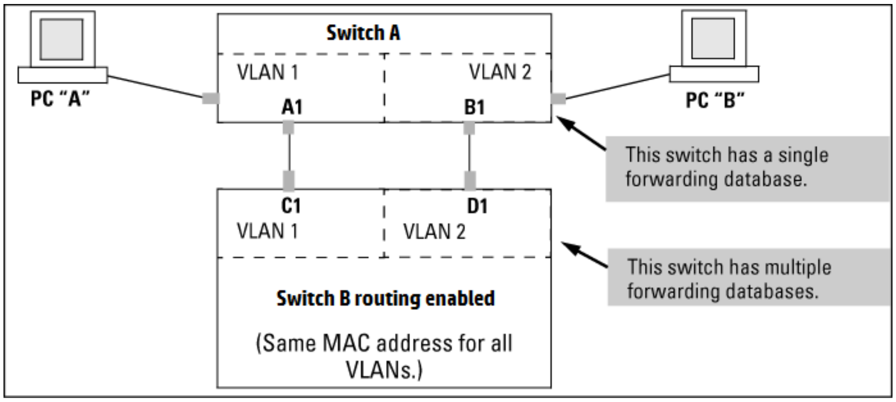

Switch A
VLAN 1
VLAN 2
A1
B1
PC "A"
PC "B"
C1
D1
VLAN 1
VLAN 2
Switch B routing enabled
(Same MAC address for all VLANs.)
This switch has a single forwarding database.
This switch has multiple forwarding databases.

Action/solution

Reconfigure the switches in the configuration.

Procedure

1. Use only one cable or port trunk between single-forwarding and multiple-forwarding database devices.
2. Configure the link with multiple, tagged VLANs.
3. To increase network bandwidth of the connection between devices, use a trunk of multiple physical links.

Following these rules, the switch forwarding database always lists the switch MAC address on port A1 and the switch will send traffic to either VLAN on the switch.

Figure 7: *Solution for single-forwarding to multiple-forwarding database devices in a multiple VLAN environment*

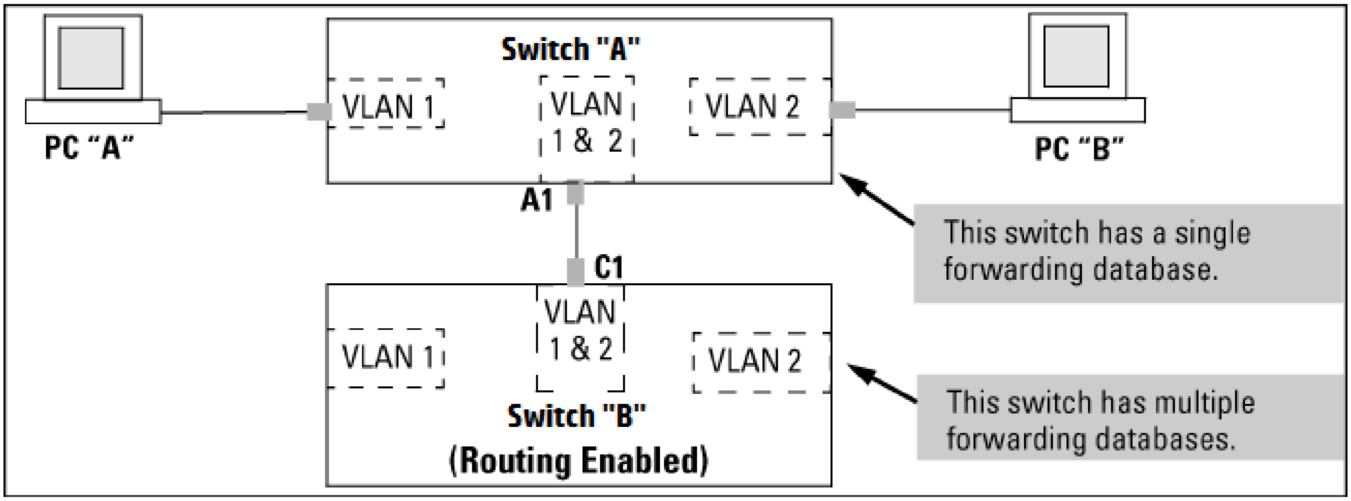

Connecting the Switch to another switch with a multiple forwarding database (Example)

Use one or both of the following connection options:

- A separate port or port trunk interface for each VLAN. This results in a forwarding database having multiple instances of the same MAC address with different VLAN IDs and port numbers. See **Forwarding database content**. The fact that the switches covered by this guide use the same MAC address on all VLAN interfaces causes no problems.
- The same port or port trunk interface for multiple (tagged) VLANs. This results in a forwarding database having multiple instances of the same MAC address with different VLAN IDs, but the same port number.

Allowing multiple entries of the same MAC address on different VLANs enables topologies such as the following:

Figure 8: *Topology for devices with multiple forwarding databases in a multiple VLAN environment*

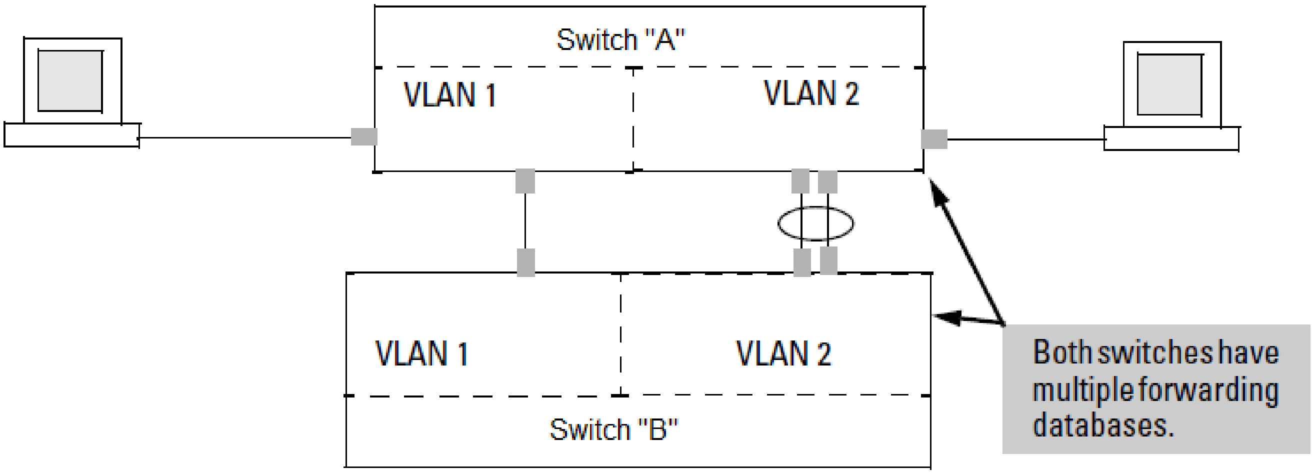

Configuring VLANs

The CLI configures and displays port-based and protocol-based VLANs.

In the factory default state, the switch is enabled for up to 256 VLANs, all ports belong to the default primary VLAN and are in the same broadcast/multicast domain. You can reconfigure the switch to support more VLANs . The maximum VLANs allowed varies according to the switch series.

The number of VLANs allowed on a switch

The factory default number of VLANs is 256.

You can reconfigure the switch to support more VLANs using the `max-vlans` command or the GUI. The maximum VLANs allowed varies according to the switch series. The maximum VLAN values for the switch documented in this guide are as follows:

Attribute	MAX Number of VLANs
2530 Switch Series; YA/YB code, 2540 Switch Series; YC code	
VLAN	512

Table Continued

Attribute	MAX Number of VLANs
IP VLAN	512 total with up to: • 512 IPv4 • 512 IPv6
static routes	256 total

The maximum VIDs is 4094.

Per-port static VLAN configuration options example

This example shows the options available to assign individual ports to a static VLAN.

GVRP, if configured, affects these options and the VLAN behavior on the switch.

Figure 9: *Comparing per-port VLAN options with and without GVRP*

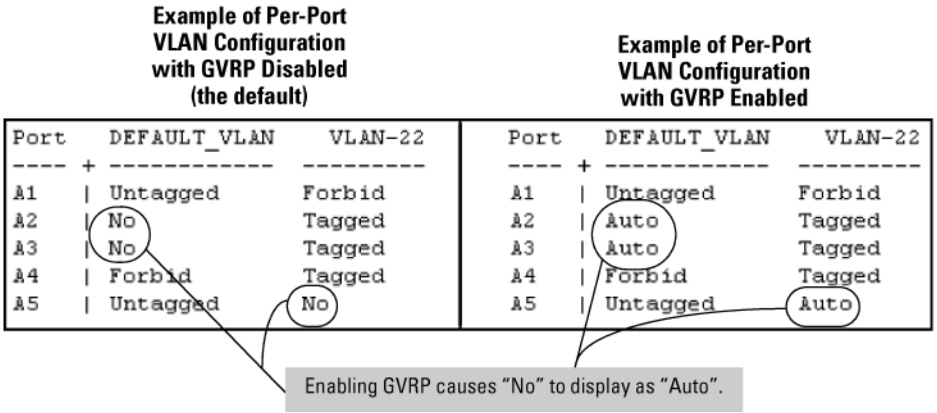

Example of Per-Port VLAN Configuration with GVRP Disabled (the default)

```
Port    DEFAULT_VLAN     VLAN-22
---- + ------------    ---------
A1   | Untagged        Forbid
A2   | No              Tagged
A3   | No              Tagged
A4   | Forbid          Tagged
A5   | Untagged        No
```

Example of Per-Port VLAN Configuration with GVRP Enabled

```
Port    DEFAULT_VLAN     VLAN-22
---- + ------------    ---------
A1   | Untagged        Forbid
A2   | Auto            Tagged
A3   | Auto            Tagged
A4   | Forbid          Tagged
A5   | Untagged        Auto
```

Enabling GVRP causes “No” to display as “Auto”.

Table 4: *Per-port VLAN configuration options*

Parameter	Effect on port participation in designated VLAN
`Tagged`	Allows the port to join multiple VLANs.
`Untagged`	• Allows VLAN connection to a device that is configured for an untagged VLAN instead of a tagged VLAN. • A port can be an untagged member of only one port-based VLAN. • A port can be an untagged member of only one protocol-based VLAN for any given protocol type. For example, if the switch is configured with the default VLAN plus three protocol-based VLANs that include IPX, then port 1 can be an untagged member of the default VLAN and one of the protocol-based VLANs.

Table Continued

Parameter	Effect on port participation in designated VLAN
No or Auto	No: When the switch is not GVRP-enabled; prevents the port from joining that VLAN. Auto: When GVRP is enabled on the switch; it allows the port to dynamically join any advertised VLAN that has the same VID.
Forbid	Prevents the port from joining the VLAN, even if GVRP is enabled on the switch.

Configuring port-based VLAN parameters

NOTE: The CLI configures and displays both port-based and protocol-based VLANs (see **Using the CLI to configure port-based and protocol-based VLAN parameters** on page 32.

In the factory default state, the switch is enabled for up to 256 VLANs, all ports belong to the default primary VLAN and are in the same broadcast/multicast domain. The default VLAN is also the default Primary VLAN; see **The primary VLAN** on page 53. In addition to the default VLAN, you can configure additional static VLANs by adding new VLAN names and VIDs, and then assigning one or more ports to each VLAN. (The maximum of VLANs includes the default VLAN, all additional static VLANs you configure, and any dynamic VLANs the switch creates if you enable GVRP; see **GVRP** on page 62.) Each port can be assigned to multiple VLANs by using VLAN tagging; see **VLAN tagging rules** on page 21.)

Using the CLI to configure port-based and protocol-based VLAN parameters

In the factory default state, all ports on the switch belong to the port-based default VLAN (DEFAULT_VLAN; VID=1) and are in the same broadcast/multicast domain.

The default VLAN is also the Primary VLAN.

You can configure additional static VLANs by adding new VLAN names and then assigning one or more ports to each VLAN.

The maximum VLANs accepted by the switch varies according to the switch series. VIDs numbered up to 4094 are allowed. This must include the default VLAN and any dynamic VLANs the switch creates if you enable GVRP (see **GVRP** on page 62).

NOTE: Each port can be assigned to multiple VLANs by using VLAN tagging. See **VLAN tagging rules** on page 21.

Creating a new static VLAN (port-based or protocol-based) (CLI)

The `vlan <vid>` command operates in the global configuration context to configure a static VLAN and/or take the CLI to a specified VLAN's context.

Syntax:

```
vlan <vid> | <ascii-name-string>
```

```
no vlan <vid>
```

If *<vid>* does not exist in the switch, this command creates a port-based VLAN with the specified *<vid>*

If the command does not include options, the CLI, moves to the newly created VLAN context.

If an optional name is not specified, the switch assigns a name in the default format `VLAN` **n**, where `n` is the *`<vid>`* assigned to the VLAN.

If the VLAN exists and you enter either the *`<vid>`* or the *`<ascii-name-string>`* ,the CLI moves to the specified VLAN's context.

The `no` form of the command deletes the VLAN as follows:

If one or more ports belong only to the VLAN to be deleted, the CLI notifies you that these ports will be moved to the default VLAN and prompts you to continue the deletion. For member ports that also belong to another VLAN, there is no **move** prompt.

```
protocol [ipx | ipv4 | ipv6 | arp | appletalk | sna | netbeui]
```

Configures a static, protocol VLAN of the specified type.

If multiple protocols are configured in the VLAN, the `no` form removes the specified protocol

If a protocol VLAN is configured with only one protocol type and you use the `no` form of this command to remove that protocol, the switch changes the protocol VLAN to a port-based VLAN (if the VLAN does not have an untagged member port).

If an untagged member port exists on the protocol VLAN, you must either convert the port to a tagged member or remove the port from the VLAN before removing the last protocol type from the VLAN.

NOTE: If you create an IPv4 protocol VLAN, you must assign the ARP protocol option to it to provide IP address resolution. Otherwise, IP packets are not deliverable. A Caution message appears in the CLI if you configure IPv4 in a protocol VLAN that does not already include the ARP protocol option. The same message appears if you add or delete another protocol in the same VLAN.

```
name <ascii-name-string>
```

When included in a `vlan` command to create a new static VLAN, this command specifies a non-default VLAN name. Also used to change the current name of an existing VLAN.

NOTE: Avoid spaces and the following characters in the *`<ascii-name-string>`* entry: `@`, `#:`, `$`, `^`, `&`, `*`, `(` and`)`. To include a blank space in a VLAN name, enclose the name in single or double quotes.

```
voice
```

Designates a VLAN for VoIP use. For more on this topic, see **Using voice VLANs** on page 52.

NOTE: You can use these options from the configuration level by beginning the command with `vlan` *`<vid>`* , or from the context level of the specific VLAN by just entering the command option.

Creating a new port-based static VLAN

The following example shows how to create a new port-based, static VLAN with a VID of 100 using the following steps:

1. To create the new VLAN, type the `vlan 100` command.
2. To show the VLANs currently configured in the switch, type the `show vlans` command.

If the Management VLAN field (`Primary VLAN : DEFAULT_VLAN Management VLAN` shown in the display information below) is empty, a Secure Management VLAN is not configured in the switch. For more information on configuring a secure management VLAN, see **The secure Management VLAN** on page 54.

```
switch(config)# vlan 100
switch(config)# show vlans

 Status and Counters - VLAN Information
 Maximum VLANs to support : 16
 Primary VLAN : DEFAULT_VLAN
 Management VLAN :

 VLAN ID Name                 Status       Voice Jumbo
 ------- -------------------- ------------ ----- -----
 1       DEFAULT_VLAN         Port-based   No    No
 100     VLAN100              Port-based   No    No
```

Changing the VLAN context level

To go to a different VLAN context level, such as to the default VLAN:

```
switch(vlan-100)# vlan default_vlan
switch(vlan-1)# _
```

Configuring or changing static VLAN per-port settings (CLI)

Syntax:

`vlan <vid>`

`no vlan <vid>`

This command, used with the options listed below, changes the name of an existing static VLAN and the per-port VLAN membership settings.

NOTE: You can use these options from the configuration level by beginning the command with `vlan <vid>`, or from the context level of the specific VLAN by just entering the command option.

`tagged <port-list>`

Configures the indicated port as Tagged for the specified VLAN. The `no` version sets the port to either **No** or (if GVRP is enabled) to **Auto**.

`untagged <port-list>`

Configures the indicated port as Untagged for the specified VLAN. The `no` version sets the port to either **No** or (if GVRP is enabled) to **Auto**.

`forbid <port-list>`

Used in port-based VLANs, configures `<port-list>` as forbidden to become a member of the specified VLAN, as well as other actions. Does not operate with option not allowed protocol VLANs. The `no` version sets the port to either `No` or (if GVRP is enabled) to `Auto`. See **GVRP** on page 62.

`auto <port-list>`

Available if GVRP is enabled on the switch. Returns the per-port settings for the specified VLAN to `Auto` operation. `Auto` is the default per-port setting for a static VLAN if GVRP is running on the switch. For information on dynamic VLAN and GVRP operation, see **GVRP** on page 62.

Changing the VLAN name and set ports to tagged

Suppose that there is a VLAN named VLAN100 with a VID of 100 and all ports are set to **No** for this VLAN. To change the VLAN name to `Blue_Team` and set ports A1 - A5 to Tagged, use the following commands:

```
switch(config)# vlan 100 name Blue_Team
switch(config)# vlan 100 tagged a1-a5
```

Moving the context level

To move to the `vlan 100` context level and execute the same commands:

```
switch(config)# vlan 100
switch(vlan-100)# name Blue_Team
switch(vlan-100)# tagged a1-a5
```

Changing tagged ports

Similarly, to change the tagged ports in the above examples to `No` (or `Auto`, if GVRP is enabled), use either of the following commands.

At the global config level, use:

```
switch(config)# no vlan 100 tagged a1-a5
```

or

At the VLAN 100 context level, use:

```
switch(vlan-100)# no tagged a1-a5
```

NOTE: You cannot use these commands with dynamic VLANs. Attempting to do so displays the message `VLAN already exists` with no change.

Converting a dynamic VLAN to a static VLAN (CLI)

Syntax:

`static-vlan <vlan-id>`

Converts a dynamic, port-based VLAN membership to static, port-based VLAN membership (allows port-based VLANs only).

For this command,*<vlan-id>* refers to the VID of the dynamic VLAN membership. Use `show vlan` to help identify the VID.

This command requires that GVRP is running on the switch and a port is currently a dynamic member of the selected VLAN.

After you convert a dynamic VLAN to static, you must configure the switch's per-port participation in the VLAN in the same way that you would for any static VLAN. For GVRP and dynamic VLAN operation, see **GVRP** on page 62.

Converting a dynamic VLAN to a port-based static VLAN

Suppose a dynamic VLAN with a VID of 125 exists on the switch. The following command converts the VLAN to a port-based, static VLAN:

```
switch(config)# static-vlan 125
```

Deleting a static VLAN (CLI)

Syntax:

```
vlan <vid>

no vlan <vid>
```

CAUTION: Before deleting a static VLAN, reassign all ports in the VLAN to another VLAN.

Deleting a static VLAN

If ports B1-B5 belong to both VLAN 2 and VLAN 3 and ports B6-B10 belong to VLAN 3, deleting VLAN 3 causes the CLI to prompt you to approve moving ports B6 - B10 to VLAN 1 (the default VLAN). (Ports B1-B5 are not moved because they still belong to another VLAN.)

```
switch(config)# no vlan 3
The following ports will be moved to the default VLAN:
B6-B10
Do you want to continue?
[y/n] Y
switch(config)#
```

Deleting multiple VLANs

The `interface` command enables you to add or delete interfaces from multiple tagged or untagged VLANs or SVLANs using a single command. Interfaces can be added or deleted for up to 256 VLANs at a time. If more than 256 VLANs are specified, an error is displayed. The `forbid` option prevents an interface from becoming a member of the specified VLANs or SVLANs when used with GVRP.

Syntax

```
interface <port-list> <tagged | untagged | forbid> <vlan | svlan <vlan-id-list>>
```

```
no interface <port-list> <tagged | untagged | forbid> <vlan | svlan <vlan-id-list>>
```

The specified interfaces are added to existing VLANs or SVLANs. If a VLAN or SVLAN does not exist, an error message displays.

The `no` form of the command removes the specified interfaces from the specified VLANs or SVLANs.

The `forbid` option prevents an interface from becoming a member of the specified VLANs or SVLANs. It is executed in interface context.

Removing an interface from several VLANs

The `vlan-id-list` includes a comma-separated list of VLAN IDs and/or VLAN ID ranges.

To remove interface 1 from VLANs 1, 3, 5, 6, 7, 8, 9, 10

```
switch(config)# no interface 1,6,7-10 tagged vlan 1,3,5-10
```

To specify that an interface cannot become a member of VLANs 4 and 5

```
switch(config)# interface 2 forbid vlan 4-5
```

Using IP enable/disable for all VLANs

You can administratively disable the IP address on specified VLANs with static IP addresses without removing the Layer 3 configuration. The switch can be pre-configured as a backup router, then quickly transition from backup to active by re-enabling Layer 3 routing on one or more VLANs. While the switch is in “backup” mode, it will still be performing Layer 2 switching.

A MIB object will be toggled to make Layer 3 routing active or inactive on a VLAN.

Interaction with other features

This feature affects management access to the switch as follows:

- IP—SNMP, Telnet, SSH, HTTP, TFTP, SCP, SFTP
- Routing—RIP, OSPF, PIM, VRRP

When the `disable layer3` command is configured on a VLAN, the behavior is as if no IP address were configured for that VLAN. There is no other change in behavior.

Syntax:

```
disable layer3 vlan <vid> <vid range>
```

```
no disable layer3 vlan <vid> <vid range>
```

In config context, turns off Layer 3 routing for the specified VLAN or VLANs. When executed in vlan context, turns off Layer 3 routing for that VLAN.

The `no` form turns on Layer 3 routing for the specified VLAN or VLANs.

The `show ip` command displays `disabled` in the IP Config column if Layer 3 has been disabled, or if the VLAN has no IP configuration. You can tell which is the case by viewing the remaining columns; if there is no IP configuration, the remaining columns are blank.

Displaying a VLAN disabled for Layer 3

```
switch(config)# show ip

 Internet (IP) Service

  IP Routing : Disabled

  Default Gateway : 172.22.16.1
  Default TTL     : 64
  Arp Age         : 20
  Domain Suffix   :
  DNS server      :

                                                                 Proxy ARP
  VLAN              | IP Config  IP Address      Subnet Mask      Std   Local
  ----------------- + ---------- --------------- --------------- ----  -----
  DEFAULT_VLAN      | DHCP/Bootp 172.22.18.100   255.255.248.0   No    No
  VLAN3             | Disabled   172.17.17.17    255.255.255.0   No    No
  VLAN6             | Disabled
  VLAN7             | Manual     10.7.7.1        255.255.255.0   No    No
```

For IPv6, the `Layer 3 Status` field displays the status of Layer 3 on that VLAN.

Displaying IPv6 Layer 3 status for a VLAN

```
switch(config)# show ipv6

 Internet (IPv6) Service

  IPv6 Routing    : Disabled
  Default Gateway :
  ND DAD          : Enabled
  DAD Attempts    : 3

  Vlan Name       : DEFAULT_VLAN
  IPv6 Status     : Disabled
  Layer 3 Status  : Enabled

  Vlan Name       : layer3_off_vlan
  IPv6 Status     : Disabled
  Layer 3 Status  : Disabled

  Address    |                                             Address
  Origin     | IPv6 Address/Prefix Length                  Status
---------- + ------------------------------------------- -----------
  manual     | abcd::1234/32                               tentative
  autoconfig | fe80::218:71ff:febd:ee00/64                 tentative
```

Interactions with DHCP

Disabling Layer 3 functionality and DHCP are mutually exclusive, with DHCP taking precedence over disable layer3 on a VLAN. The following interactions occur:

- If the `disable layer3` command is executed when DHCP is already configured, no disabling of the VLAN occurs. This error message displays: “Layer 3 cannot be disabled on a VLAN that has DHCP enabled.”
- From the CLI: If `disable layer3` is configured already and an attempt is made to configure DHCP, DHCP takes precedence and will be set. The warning message displays: “Layer 3 has also been enabled on this VLAN since it is required for DHCP.”
- From the CLI: When disabling a range of VLAN IDs, this warning message displays: “Layer 3 will not be disabled for any LANs that have DHCP enabled.”
- From SNMP: If the `disable layer3` command is executed when DHCP is already configured, no disabling of the VLAN occurs. An INCONSISTENT_VALUE error is returned.
- From SNMP: If `disable layer3` is configured already and an attempt is made to configure DHCP, DHCP takes precedence and will be set.

Changing the Primary VLAN (CLI)

For more information on Primary VLANs, see **The primary VLAN** on page 53.

To change the Primary VLAN (CLI), use the following command:

```
primary-vlan vid <ascii-name-string>
```

In the default VLAN configuration, the port-based default VLAN (`DEFAULT_VLAN`) is the Primary VLAN. This command reassigns the Primary VLAN function to an existing, port-based, static VLAN. The switch cannot reassign the Primary VLAN function to a protocol VLAN.

If you reassign the Primary VLAN to a non-default VLAN, to delete the Primary VLAN from the switch, you must assign the Primary VLAN to another port-based static VLAN.

To identify the current Primary VLAN and list the available VLANs and their respective VIDs, use `show vlans`.

Reassigning, renaming and displaying the VLAN command sequence

The following example shows how to reassign the Primary VLAN to VLAN 22 (first command line), rename the VLAN **`22-Primary`** (second command line) and then display the result (third command line):

```
switch(config)# primary-vlan 22
switch(config)# vlan 22 name 22-Primary
switch(config)# show vlans

Status and Counters - VLAN Information
Maximum VLANs to support : 8
Primary VLAN : 22-Primary
Management VLAN :

VLAN ID Name                 Status       Voice Jumbo
------- -------------------- ------------ ----- -----
1       DEFAULT_VLAN         Static       No    No
22      22-Primary           Static       No    No
```

Configuring a secure Management VLAN (CLI)

Preparation

Procedure

1. Determine a VID and VLAN name suitable for your Management VLAN.
2. Plan your topology to use switches that support Management VLANs. See **The secure Management VLAN** on page 54.
3. Include only the following ports:
   a. Ports to which you will connect authorized management stations, such as Port A7 in the "Management VLAN control in a LAN" example in **The secure Management VLAN**.
   b. Ports on one switch that you will use to extend the Management VLAN to ports on other switches, such as ports A1 in the "Management VLAN control in a LAN" example in **The secure Management VLAN**.
4. Half-duplex repeaters dedicated to connecting management stations to the Management VLAN can also be included in this topology. Any device connected to a half-duplex repeater in the Management VLAN will also have Management VLAN access.
5. Configure the Management VLAN on the selected switch ports.
6. Test the Management VLAN from all of the management stations authorized to use it, including any SNMP-based network management stations. Also test any Management VLAN links between switches.

NOTE: If you configure a Management VLAN on a switch using a Telnet connection through a port not in the Management VLAN, you will lose management contact with the switch if you log off your Telnet connection or execute `write memory` and `reboot` the switch.

Configuring an existing VLAN as the Management VLAN (CLI)

Syntax:

```
management-vlan <vlan-id> | <vlan-name>

no management-vlan <vlan-id> | <vlan-name>
```

Configures an existing VLAN as the Management VLAN.

The `no` form disables the Management VLAN and returns the switch to its default management operation.

Default: Disabled. In this case, the VLAN returns to standard VLAN operation.

Switch configuration

You have configured a VLAN named `My_VLAN` with a VID of 100 and want to configure the switch to do the following:

- Use `My_VLAN` as a Management VLAN (tagged, in this case) to connect port A1 on switch "A" to a management station. The management station includes a network interface card with 802.1Q tagged VLAN capability.
- Use port A2 to extend the Management VLAN to port B1 which is already configured as a tagged member of `My_VLAN`, on an adjacent switch that supports the Management VLAN feature.

```
switch(config)# management-vlan 100
switch(config)# vlan 100 tagged a1
switch(config)# vlan 100 tagged a2
```

Configuration Example

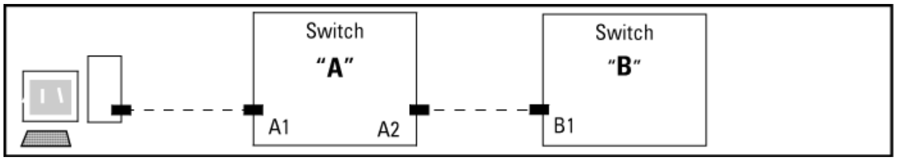

Obtaining an IP address using DHCP (CLI)

Use DHCP to obtain an IPv4 address for your Management VLAN or a client on that VLAN. The following examples illustrate when an IP address will be received from the DHCP server.

DHCP server on a Management VLAN

If Blue_VLAN is configured as the Management VLAN and the DHCP server is also on Blue_VLAN, Blue_VLAN receives an IP address. Because DHCP Relay does not forward onto or off the Management VLAN, devices on Red_VLAN cannot get an IP address from the DHCP server on Blue_VLAN (Management VLAN) and Red_VLAN does not receive an IP address.

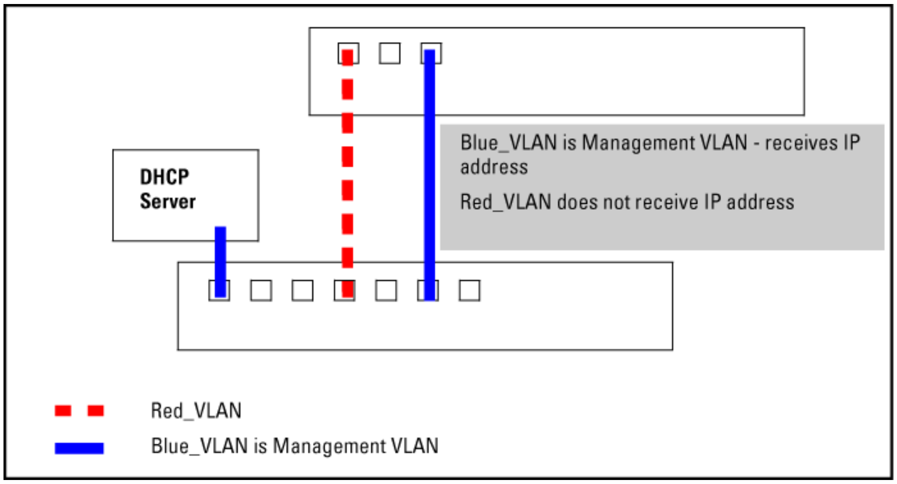
DHCP
Server
Blue_VLAN is Management VLAN - receives IP address
Red_VLAN does not receive IP address
Red_VLAN
Blue_VLAN is Management VLAN

DHCP server on a different VLAN from the Management VLAN

If Red_VLAN is configured as the Management VLAN and the DHCP server is on Blue_VLAN, Blue_VLAN receives an IP address but Red_VLAN does not.

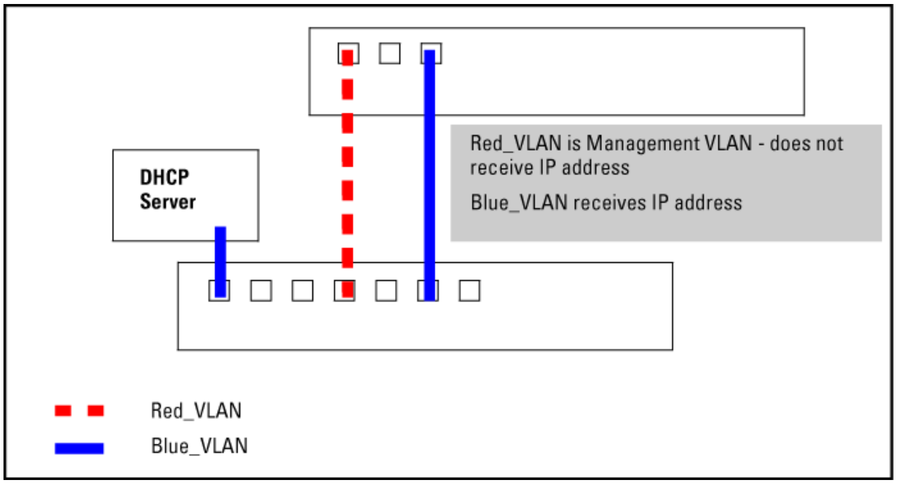

Red_VLAN is Management VLAN - does not receive IP address

Blue_VLAN receives IP address

No Management VLANs configured

If no Management VLAN is configured, both Blue_VLAN and Red_VLAN receive IP addresses.

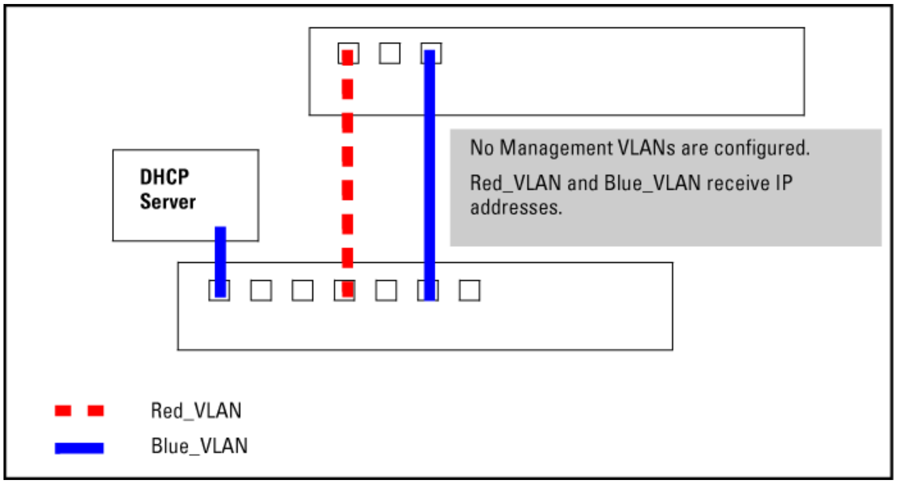

No Management VLANs are configured.

Red_VLAN and Blue_VLAN receive IP addresses.

A client on a different Management VLAN from the DHCP server

If Red_VLAN is configured as the Management VLAN and the client is on Red_VLAN, but the DHCP server is on Blue_VLAN, the client will not receive an IP address.

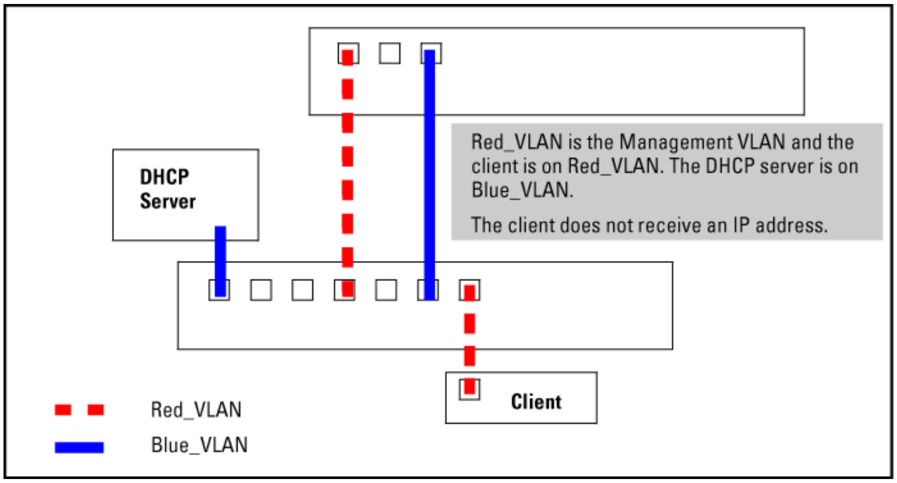

Red_VLAN is the Management VLAN and the client is on Red_VLAN. The DHCP server is on Blue_VLAN.

The client does not receive an IP address.

A DHCP server and client on the Management VLAN

If Blue_VLAN is configured as the Management VLAN, the client is on Blue_VLAN and the DHCP server is on Blue_VLAN, the client receives an IP address.

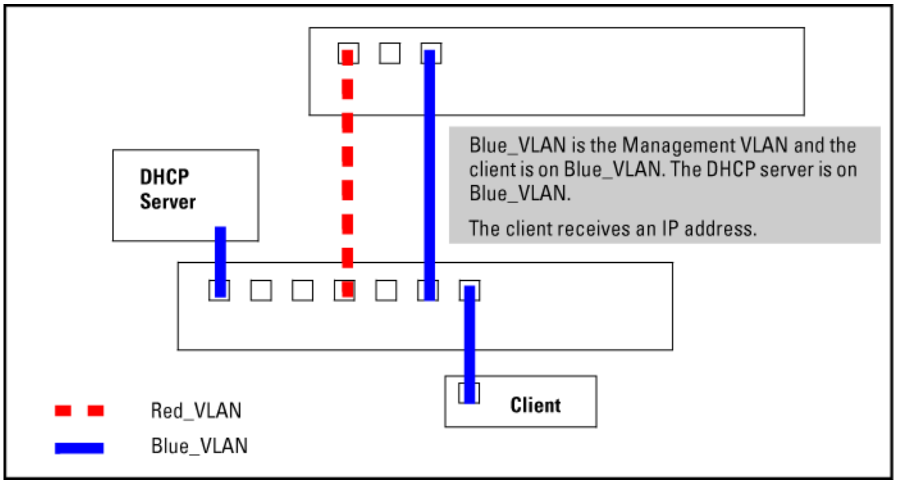
DHCP Server
Blue_VLAN is the Management VLAN and the client is on Blue_VLAN. The DHCP server is on Blue_VLAN.
The client receives an IP address.
Client
Red_VLAN
Blue_VLAN

Obtaining the IP address for a host that is on a different VLAN than the DHCP server

In the following example, the host is on VLAN 20 and is connected on port number 2 of the switch. The DHCP server, however, is in VLAN 10 and is connected on port 10 of the switch.

Obtaining the IP address for a host that is on a different VLAN than the DHCP server

```
switch(config)# vlan 10
name "VLAN 10"
untagged 10
ip address 10.1.1.2 255.255.255.0
exit
vlan 20
name "VLAN 20"
untagged 2
ip address 100.99.1.1 255.255.255.0
ip helper-address 10.1.1.1
exit
```

Disabling the Management feature (CLI)

You can disable the Secure Management feature without deleting the VLAN.

Disabling the secure management feature

The following commands disable the Secure Management feature in the above example:

```
switch(config)# no management-vlan 100
switch(config)# no management-vlan my_vlan
```

For more information, see **The secure Management VLAN** on page 54.

Changing the number of VLANs allowed on the switch (CLI)

Syntax:

```
max-vlans<max number of vlans>
```

Use this command to specify the maximum number of VLANs allowed on the switch. The minimum value is 16. The maximum value varies according to the switch series.

For the 2530 switch series you can enter a `max-vlans` value of between 16–512.

The total number of allowed IP VLANs (IPv6 + IPv4) is 512.

If GVRP is enabled, this setting includes any dynamic VLANs on the switch. As part of implementing a new setting, you must execute a `write memory` command to save the new value to the startup-config file and then reboot the switch.

NOTE: If multiple VLANs exist on the switch, you cannot reset the maximum number of VLANs to a value smaller than the current number of VLANs.

The following example shows the command sequence for changing the number of VLANs allowed to 20. You can execute the commands to `write memory` and `boot` at another time.

Example of changing the number of allowed VLANs

```
switch(config)# max-vlans 20
This command will take effect after saving the configuration
and rebooting the system.
switch(config)# write memory
switch(config)# boot
This will reboot the system from the primary image, do you want to continue [y/n]? Y
```

Error Messages

An error message will be displayed, if you set the `max-vlans` value to a number that exceeds the allowable value for the switch series.

If you set the `max-vlans` and later try to downgrade to an earlier version of the switch software that does not allow that number of `max-vlans`, successful downgrade may be prevented.

Displaying a switch VLAN configuration

The `show vlans` command lists the VLANs currently running in the switch, with VID, VLAN name, and VLAN status. Dynamic VLANs appear only if the switch is running with GVRP enabled and one or more ports has dynamically joined an advertised VLAN. In the default configuration, GVRP is disabled.

Syntax:

```
show vlans
```

The following describes the fields displayed with this command (see the example output):

Maximum VLANs to support

Shows the number of VLANs the switch is currently configured to support.

Primary VLAN

See **The primary VLAN** on page 53.

Management VLAN

See **The secure Management VLAN** on page 54.

802.1Q VLAN ID

The VLAN identification number, or VID.

Name

The default or specified name assigned to the VLAN. For a static VLAN, the default name consists of `VLAN-x` where `x` matches the VID assigned to that VLAN. For a dynamic VLAN, the name consists of `GVRP_x` where `x` matches the applicable VID.

Status

Port-Based

Port-Based, static VLAN

Protocol

Protocol-Based, static VLAN

Dynamic

Port-Based, temporary VLAN learned through GVRP

Voice

Indicates whether a port-based VLAN is configured as a voice VLAN. See **Using voice VLANs** on page 52.

Jumbo

Indicates whether a VLAN is configured for Jumbo packets. For more on jumbos, see "Port Traffic Controls" in the management and configuration guide for your switch.

This example shows the listing from the `show vlans` command. When GVRP is disabled (the default), Dynamic VLANs do not exist on the switch and do not appear in this listing. For more information, see **GVRP** on page 62.

Displaying VLAN listing with GVRP enabled

```
switch# show vlans

 Status and Counters - VLAN Information

  Maximum VLANs to support : 256
  Primary VLAN : DEFAULT_VLAN
  Management VLAN :

  VLAN ID Name                 | Status     Voice Jumbo
  ------- -------------------- + ---------- ----- -----
  1       DEFAULT_VLAN         | Port-based No    No
  10      VLAN_10              | Port-based Yes   Yes
  15      VLAN_15              | Port-based No    No
  20      VLAN_20              | Protocol   No    No
  33      VLAN_33              | Dynamic    No    No
```

Viewing the VLAN membership of one or more ports (CLI)

Syntax:

```
show vlan ports <port-list> [detail]
```

Displays VLAN information for an individual port or a group of ports, either cumulatively or on a detailed per-port basis.

`port-list`

Specifies a single port number or a range of ports (for example, `a1-a16`), or `all` for which to display information.

`detail`

Displays detailed VLAN membership information on a per-port basis.

The following describes the fields displayed by the command (see example output):

Port name

The user-specified port name, if one has been assigned.

VLAN ID

The VLAN identification number, or VID.

Name

The default or specified name assigned to the VLAN. For a static VLAN, the default name consists of `VLAN-x` where `x` matches the VID assigned to that VLAN. For a dynamic VLAN, the name consists of `GVRP_x` where `x` matches the applicable VID.

Status

Port-Based

Port-Based, static VLAN.

Protocol

Protocol-Based, static VLAN.

Dynamic

Port-Based, temporary VLAN learned through GVRP.

Voice

Indicates whether a port-based VLAN is configured as a voice VLAN.

Jumbo

Indicates whether a VLAN is configured for jumbo packets. For more on jumbos, see "Port Traffic Controls" in the management and configuration guide for your switch.

Mode

Indicates whether a VLAN is tagged or untagged.

Displaying VLAN ports (cumulative listing)

```
switch(config)#show vlan ports a1-a24

Status and Counters - VLAN Information - for ports A1-A24

VLAN ID Name                 | Status     Voice Jumbo
------- -------------------- + ---------- ----- -----
1       DEFAULT_VLAN         | Port-based No    No
10      VLAN_10              | Port-based Yes   No
15      VLAN_15              | Protocol   No    No
```

Displaying VLAN ports (detailed listing)

```
switch(config)#show vlan ports a1-a3 detail
```

```
Status and Counters - VLAN Information - for ports A1

VLAN ID Name                 | Status     Voice Jumbo Mode
------- -------------------- + ---------- ----- ----- --------
1       DEFAULT_VLAN         | Port-based No    No    Untagged
10      VLAN_10              | Port-based Yes   No    Tagged

Status and Counters - VLAN Information - for ports A2

VLAN ID Name                 | Status     Voice Jumbo Mode
------- -------------------- + ---------- ----- ----- --------
1       DEFAULT_VLAN         | Port-based No    No    Untagged
20      VLAN_20              | Protocol   No    No    Untagged

Status and Counters - VLAN Information - for ports A3

VLAN ID Name                 | Status     Voice Jumbo Mode
------- -------------------- + ---------- ----- ----- --------
1       DEFAULT_VLAN         | Port-based No    No    Untagged
33      VLAN_33              | Port-based No    No    Tagged
```

Viewing the configuration for a particular VLAN (CLI)

Syntax:

```
show vlans <vlan-id>
```

Uses the VID to identify and display the data for a specific static or dynamic VLAN.

The following describes the fields displayed with this command (see example output):

802.1Q VLAN ID

The VLAN identification number, or VID.

Name

The default or specified name assigned to the VLAN. For a static VLAN, the default name consists of `VLAN-x` where `x` matches the VID assigned to that VLAN. For a dynamic VLAN, the name consists of `GVRP_x` where `x` matches the applicable VID.

Status

Port-Based

Port-Based, static VLAN.

Protocol

Protocol-Based, static VLAN

Dynamic

Port-Based, temporary VLAN learned through GVRP. See **GVRP** on page 62.

Voice

Indicates whether a port-based VLAN is configured as a voice VLAN. See **Using voice VLANs** on page 52.

Jumbo

Indicates whether a VLAN is configured for Jumbo packets. For more on jumbos, see "Port Traffic Controls" in the management and configuration guide for your switch.

Port Information

Lists the ports configured as members of the VLAN.

DEFAULT

Shows whether a port is a tagged or untagged member of the listed VLAN.

Unknown VLAN

Shows whether the port can become a dynamic member of an unknown VLAN for which it receives an advertisement. GVRP must be enabled to allow dynamic joining to occur.

Status

Shows whether the port is participating in an active link.

Displaying information for a specific static VLAN

```
switch(config)#show vlans 22

Status and Counters - VLAN Information - VLAN 22

 VLAN ID : 22
 Name : VLAN22
 Status : Port-based
 Voice : Yes
 Jumbo : No

 Port Information Mode     Unknown VLAN Status
 ---------------- -------- ------------ ----------
 12               Untagged Learn        Up
 13               Untagged Learn        Up
 14               Untagged Learn        Up
 15               Untagged Learn        Down
 16               Untagged Learn        Up
 17               Untagged Learn        Up
 18               Untagged Learn        Up
```

Displaying information for a specific dynamic VLAN

The following example shows the information displayed for a specific dynamic VLAN. The `show vlans` command lists this data when GVRP is enabled and at least one port on the switch has dynamically joined the designated VLAN.

```
switch(config)# show vlans 22

 Status and Counters - VLAN Information - VLAN 22

  VLAN ID : 33
  Name : GVRP_33
  Status : Dynamic
  Voice : No
  Jumbo : No

  Port Information Mode     Unknown VLAN Status
  ---------------- -------- ------------ ----------
  6                Auto     Learn        Up
```

Customizing the show VLANs output (CLI)

Syntax

```
show vlans custom [port <port-list>] <column-list>
```

Specifies the order you want information to display for the `show vlans` command. Displays information for one port or a range of ports. If `<port-list>` is not specified, all ports display.

Fields that can be included in the customized display:

Field	Display	Example	Default width
id	VLAN id	5	6
name	VLAN name	Vlan55	32
status	Status	Port-based	10
voice	Voice enabled	No	5
jumbo	Jumbos enabled	No	5
ipconfig	How the IP address was configured	Manual Disabled DHCP/BootP	10
ipaddr (IPv4) ipaddr (IPv6)	The IP addresses	10.10.10.3 fe80::212:79ff:fe8d:8000	15 for IPv4 46 for IPv6
ipmask	The subnet masks	255.255.255.6/64 (prefix for IPv6 is in format "/XX")	15
proxyarp	Whether proxy ARP is configured	No	5
localproxyarp	Whether local proxy ARP is configured	No	9
state	"Up" if at least one port is up	Up	5

Customizing the VLAN display

The following example displays `id` at its default width and `name:20` allows up to 20 characters of the VLAN `name` to be displayed. The columns selected for display are separated by spaces.

If the width of the column requested is smaller than the header name of the column, the display of the header name is truncated.

```
switch(config)# show vlan custom A1-A3 id name:20 ipaddr state

Status and Counters - VLAN Information - Custom view

VLANID VLAN name            IP Addr                            State
------ -------------------- ---------------------------------  -----
1      DEFAULT_VLAN         15.255.134.74                      Up
33     Vlan33               10.10.10.01                        Up
44     Vlan44               15.255.164.13                      Up
55     Vlan55               15.255.178.2                       Down
                            15.255.178.3
                            15.255.178.4
60     Vlan60               fe80::212:79ff:fe8d:8000%vlan60    Up
```

Wrapping column headers

The total output wraps if it is longer than the terminal width; it is not truncated.

```
switch(config)# show vlan custom id
Status and Counters - VLAN Information - Custom view

 VLANID
 ------
 1
 33
 44

switch(config)# show vlan custom id:2
Status and Counters - VLAN Information - Custom view

 VL
 --
 1
 33
 44
```

Using pattern matching with the show VLANs custom command

If a pattern matching command is in a search for a field in the output of the `show vlan custom` command and it produces an error, the error message may not be visible. For example, if you enter a command with the pattern matching `include` option that contains an error (such as 'vlan' is misspelled) as in the following example, the output may be empty:

```
switch(config)# show vlans custom 1-3 name vlun include vlan1
```

Hewlett Packard Enterprise recommends that you try the `show vlans custom` command first to ensure that there is output and then enter the command again with the pattern matching option.

Creating an alias for show VLAN commands (CLI)

Create an alias for a frequently used `show vlans custom` command to avoid entering the selected columns each time you use the command.

Using a VLAN alias

```
switch(config)# alias showvlanstatus = “show vlan custom A1-A3 id name:20 status”

switch(config)# show vlan status
Status and Counters - VLAN Information - Custom view

VLANID VLAN name            Status
------ -------------------- ----------
1      DEFAULT_VLAN         Port-based
33     Vlan33               Port-based
```

Configuring a VLAN MAC address with heartbeat interval

When installing routing switches in the place of existing routers in a network configuration, you can achieve Layer 3 VLAN migration by using the `ip-recv-mac-address` command at the VLAN configuration level to:

- Configure the MAC address of the previously installed router on each VLAN interface of a routing switch.
- Optionally configure the time interval to use for sending heartbeat packets with the configured MAC address.

Syntax

```
ip-recv-mac-address <mac-address> interval <seconds>

no ip-recv-mac-address <mac-address> interval <seconds>
```

Configures a VLAN interface with the specified MAC address. Enter the `no` version of the command to remove the configured MAC address and return to the original MAC address of the switch.

Parameters

`interval <seconds>`

(Optional) Configures the time interval in seconds used between transmissions of heartbeat packets to all network devices configured on the VLAN. Valid values are from one to 255 seconds.

Default: 60 seconds.

Displaying a VLAN MAC address configuration (CLI)

Syntax:

```
show ip-recv-mac-address
```

Displaying a VLAN MAC address

```
switch# show ip-recv-mac-address

VLAN L3-Mac-Address Table
VLAN                           L3-Mac-Address                Timeout
-------------                  ------------------------      -----------
DEFAULT_VLAN                  001635-024467                  60
VLAN2                         001635-437529                  100
```

Using voice VLANs

Configuring voice VLANs separates voice traffic from data traffic and shields your voice traffic from broadcast storms.

Operating rules for voice VLANs

- You must statically configure voice VLANs. GVRP and dynamic VLANs do not support voice VLAN operation.
- Configure all ports in a voice VLAN as tagged members of the VLAN. This ensures retention of the QoS (Quality of Service) priority included in voice VLAN traffic moving through your network.
- If a telephone connected to a voice VLAN includes a data port used for connecting other networked devices (such as PCs) to the network, then you must configure the port as a tagged member of the voice VLAN and a tagged or untagged member of the data VLAN you want the other networked device to use.

Components of voice VLAN operation

- Voice VLAN: Configure one or more voice VLANs on the switch. Some reasons for having multiple voice VLANs include:
    - Employing telephones with different VLAN requirements
    - Better control of bandwidth usage
    - Segregating telephone groups used for different, exclusive purposes

  Where multiple voice VLANs exist on the switch, you can use routing to communicate between telephones on different voice VLANs.
- Tagged/Untagged VLAN Membership: If the appliances using a voice VLAN transmit tagged VLAN packets, then configure the member ports as tagged members of the VLAN. Otherwise, configure the ports as untagged members.

Voice VLAN access security

You can use port security configured on an individual port or group of ports in a voice VLAN. That is, you can allow or deny access to a phone having a particular MAC address. See chapter titled "Configuring and Monitoring Port Security" in the *Access Security Guide* for your switch.

NOTE: MAC authentication is not recommended in voice VLAN applications.

Prioritizing voice VLAN QoS (Optional)

Without configuring the switch to prioritize voice VLAN traffic, one of the following conditions applies:

- If the ports in a voice VLAN are not tagged members, the switch forwards all traffic on that VLAN at "normal" priority.
- If the ports in a voice VLAN are tagged members, then the switch forwards all traffic on that VLAN at whatever priority the traffic has when received inbound on the switch.

Using the switch's QoS VLAN-ID (VID) priority option, you can change the priority of voice VLAN traffic moving through the switch. If all port memberships on the voice VLAN are tagged, the priority level you set for voice VLAN traffic is carried to the next device. With all ports on the voice VLAN configured as tagged members, you can enforce a QoS priority policy moving through the switch and through your network.

Syntax:

```
vlan <vid> qos priority <0-7>
```

The qos priority default setting is 0 (normal), with 1 as the lowest priority and 7 as the highest priority.

If you configure a voice VLAN with a VID of 10 and want the highest priority for all traffic on this VLAN, execute the following commands:

```
switch(config)# vlan 10 qos priority 4
switch(config)# write memory
```

You also have the option of resetting the DSCP (DiffServe Codepoint) on tagged voice VLAN traffic moving through the switch. For more information, see **Quality of Service (QoS): Managing bandwidth effectively** on page 168.

If all port memberships on the voice VLAN are tagged:

- The priority level set for voice VLAN traffic is carried to the next device.
- You can enforce a QoS priority policy moving through the switch and network.

For more information, see **Using voice VLANs** on page 52.

Special VLAN types

VLAN support and the default VLAN

In the factory default configuration, VLAN support is enabled and all ports on the switch belong to the port-based, default VLAN (named DEFAULT_VLAN). This places all ports in the switch into one physical broadcast domain. In the factory-default state, the default VLAN is also the Primary VLAN.

- You can partition the switch into multiple virtual broadcast domains by configuring one or more additional VLANs and moving ports from the default VLAN to the new VLANs.
- The switch supports up to 2048 static and dynamic VLANs, with VIDs numbered up to 4094. You can change the name of the default VLAN, but not its VID, which is always 1.
- You can remove all ports from the default VLAN by placing them in another port-based VLAN, but this VLAN remains and cannot be deleted from the switch.

For details on port VLAN settings, see **Configuring or changing static VLAN per-port settings (CLI)** on page 34.

The primary VLAN

As certain features and management functions run on only one VLAN in the switch and because DHCP and Bootp can run per-VLAN, there is a need for a dedicated VLAN to manage these features and ensure that multiple instances of DHCP or Bootp on different VLANs do not result in conflicting configuration values for the switch.

The Primary VLAN is the VLAN the switch uses to run and manage these features and data. In the factory-default configuration, the switch designates the default VLAN (DEFAULT_VLAN; VID=1) as the Primary VLAN. However you can designate another static, port-based VLAN as primary.

To summarize, designating a non-default VLAN as primary means that:

- The switch reads DHCP responses on the Primary VLAN instead of on the default VLAN. This includes such DHCP-resolved parameters as the TimeP server address, Default TTL and IP addressing—including the Gateway IP address—when the switch configuration specifies DHCP as the source for these values.
- The default VLAN continues to operate as a standard VLAN you cannot delete it or change its VID.
- Any ports not specifically assigned to another VLAN will remain assigned to the Default VLAN, even if it is the Primary VLAN.

Candidates for Primary VLAN include any static, port-based VLAN currently configured on the switch.

Protocol-Based VLANs and dynamic (GVRP-learned) VLANs that have not been converted to a static VLAN cannot be the Primary VLAN. To display the current Primary VLAN, use the CLI `show vlan` command.

NOTE: If you configure a non-default VLAN as the Primary VLAN, you cannot delete that VLAN unless you first select a different VLAN to serve as primary.

If you manually configure a gateway on the switch, it ignores any gateway address received via DHCP or Bootp.

The secure Management VLAN

Configuring a secure Management VLAN creates an isolated network for managing the switches that support this feature. Access to a secure Management VLAN and the switch's management functions is available only through ports configured as members.

- Multiple ports on the switch can belong to the Management VLAN. This allows connections for multiple management stations to the Management VLAN, while allowing Management VLAN links between switches configured for the same Management VLAN.
- Only traffic from the Management VLAN can manage the switch, which means that only the workstations and PCs connected to ports belonging to the Management VLAN can manage and reconfigure the switch.

Potential security breaches in a network

This illustrates use of the Management VLAN feature to support management access by a group of management workstations.

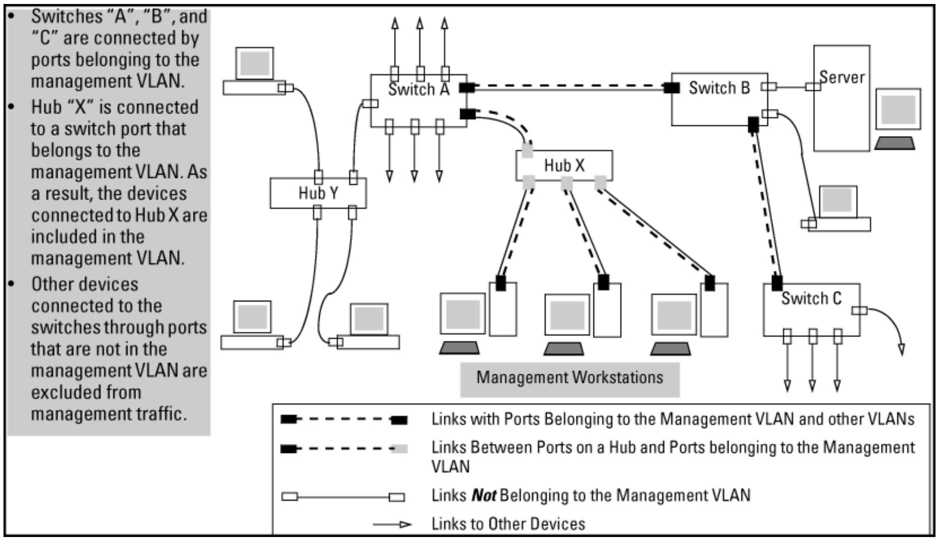
- Switches "A", "B", and "C" are connected by ports belonging to the management VLAN.
- Hub "X" is connected to a switch port that belongs to the management VLAN. As a result, the devices connected to Hub X are included in the management VLAN.
- Other devices connected to the switches through ports that are not in the management VLAN are excluded from management traffic.

Switch A
Switch B
Server
Hub X
Hub Y
Switch C
Management Workstations
Links with Ports Belonging to the Management VLAN and other VLANs
Links Between Ports on a Hub and Ports belonging to the Management VLAN
Links *Not* Belonging to the Management VLAN
Links to Other Devices

Management VLAN control in a LAN

In this example, Workstation 1 has management access to all three switches through the Management VLAN, while the PCs do not. This is because configuring a switch to recognize a Management VLAN automatically excludes attempts to send management traffic from any other VLAN.

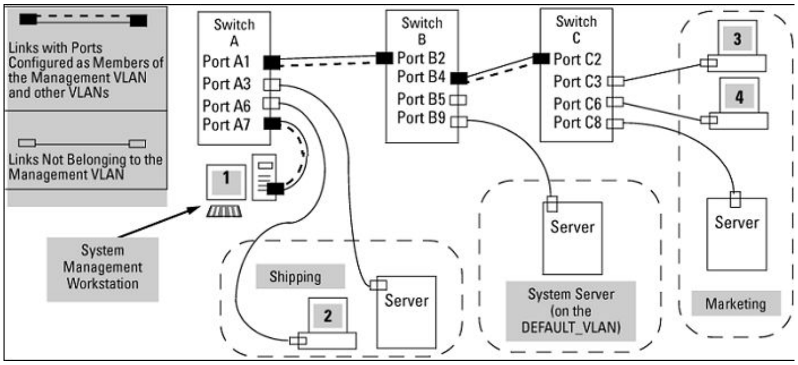
Links with Ports Configured as Members of the Management VLAN and other VLANs
Links Not Belonging to the Management VLAN
Switch A
Port A1
Port A3
Port A6
Port A7
Switch B
Port B2
Port B4
Port B5
Port B9
Switch C
Port C2
Port C3
Port C6
Port C8
3
4
1
2
System Management Workstation
Shipping
Server
Server
System Server (on the DEFAULT_VLAN)
Server
Marketing

Table 5: *VLAN membership in Management VLAN control in a LAN*

Switch	A1	A3	A6	A7	B2	B4	B5	B9	C2	C3	C6	C8
Management VLAN (VID = 7)	Y	N	N	Y	Y	Y	N	N	Y	N	N	N
Marketing VLAN (VID = 12)	N	N	N	N	N	N	N	N	N	Y	Y	Y
Shipping Dept. VLAN (VID = 20)	N	Y	Y	N	N	N	N	N	N	N	N	N
DEFAULT-VLAN (VID = 1)	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y

See **Configuring a secure Management VLAN (CLI)** on page 40 for configuration details.

Operating notes for Management VLANs

- Use only a static, port-based VLAN for the Management VLAN.
- The Management VLAN feature applies to both IPv4 and IPv6 traffic.

- The Management VLAN does not support IGMP operation.
- Routing between the Management VLAN and other VLANs is not allowed.
- If there are more than 25 VLANs configured on the switch, reboot the switch after configuring the Management VLAN.
- Only one Management VLAN can be active in the switch. If one Management VLAN VID is saved in the startup-config file and you configure a different VID in the running-config file, the switch uses the running-config version until you either use the `write-memory` command or reboot the switch.
- During a Telnet session to the switch, if you configure the Management VLAN to a VID that excludes the port through which you are connected to the switch, you will continue to have access only until you terminate the session by logging out or rebooting the switch.

NOTE: The Management VLAN feature does not control management access through a direct connection to the switch's serial port.

- During a WebAgent session, if you configure the Management VLAN to a VID that excludes the port through which you are connected to the switch, you will continue to have access only until you close the browser session or reboot the switch.
- Enabling Spanning Tree between a pair of switches where there are multiple links using separate VLANs, including the Management VLAN, will force the blocking of one or more links. This may include the link carrying the Management VLAN, which will cause loss of management access to some devices.
- Monitoring Shared Resources: The Management VLAN feature shares internal switch resources with several other features. The switch provides ample resources for all features. However, if the internal resources become fully subscribed, the Management VLAN feature cannot be configured until the necessary resources are released from other uses. For information on determining the current resource availability and usage, see the appendix titled "Monitoring Resources" in the *ArubaOS-Switch Management and Configuration Guide* for your switch.

Inadvertently blocking a Management VLAN link by implementing spanning tree

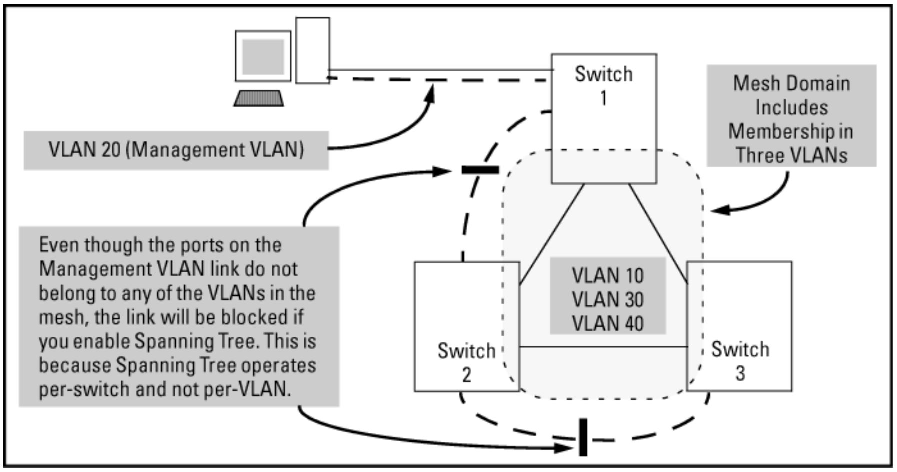
Switch 1
Mesh Domain Includes Membership in Three VLANs
VLAN 20 (Management VLAN)
Even though the ports on the Management VLAN link do not belong to any of the VLANs in the mesh, the link will be blocked if you enable Spanning Tree. This is because Spanning Tree operates per-switch and not per-VLAN.
VLAN 10
VLAN 30
VLAN 40
Switch 2
Switch 3

VLAN operating notes

DHCP/Bootp

If you are using DHCP/Bootp to acquire the switch's configuration, packet time-to-live and TimeP information, designates the VLAN on which DHCP is configured as the Primary VLAN.

NOTE:

In the factory-default configuration, the DEFAULT_VLAN is the Primary VLAN.

Per-VLAN features

IGMP and some other features operate on a per VLAN basis. This means you must configure such features separately for each VLAN in which you want them to operate.

Default VLAN

You can rename the default VLAN, but you cannot change its VID (1) or delete it from the switch.

VLAN port assignments

Any ports not specifically removed from the default VLAN remain in the DEFAULT_VLAN, regardless of other port assignments. Also, a port must always be a tagged or untagged member of at least one port-based VLAN.

Voice-Over-IP (VoIP)

VoIP operates only over static, port-based VLANs.

Multiple VLAN types configured on the same port

A port can simultaneously belong to both port-based and protocol-based VLANs.

Protocol Capacity

A protocol-based VLAN can include up to four protocol types. In protocol VLANs using the IPv4 protocol, to support normal IP network operation ARP must be one of these protocol types (to support normal IP network operation). Otherwise, IP traffic on the VLAN is disabled.

If you configure an IPv4 protocol VLAN that does not include the ARP VLAN protocol, the switch displays the following message which indicates a protocol VLAN configured with IPv4 but not ARP:

```
switch(config)# vlan 97 protocol ipv4

IPv4 assigned without ARP, this may result in undeliverable IP packets.
```

Deleting Static VLANs

A VLAN can be deleted even if there are currently ports belonging to it. The ports are moved to the default VLAN.

Adding or Deleting VLANs

Changing the number of VLANs supported on the switch, requires a reboot.

NOTE:

From the CLI, you must perform a `write memory` command before rebooting. Other VLAN configuration changes are dynamic.

Effects of VLANs on other switch features

Spanning Tree operation with VLANs

Depending on the spanning tree option configured on the switch, the spanning tree feature may operate as:

- A single instance across all ports on the switch regardless of VLAN assignments
- Multiple instances per-VLAN

For single-instance operation, if redundant physical links exist between the switch and another 802.1Q device, all but one link will be blocked, even if the redundant links are in separate VLANs. In this case, you can use port trunking to prevent Spanning Tree from unnecessarily blocking ports (and to improve overall network performance). For multiple-instance operation, physically redundant links belonging to different VLANs can remain open. For more information, see **Multiple instance spanning tree operation**.

NOTE: Spanning Tree operates differently in different devices. For example, in the (obsolete, non-802.1Q) Switch 2000 and the Switch 800T, Spanning Tree operates per-VLAN, allowing redundant physical links as long as they are in separate VLANs.

Spanning Tree operates differently in different devices

IP interfaces

There is a one-to-one relationship between a VLAN and an IP network interface. Since the VLAN is defined by a group of ports, the state (up/down) of those ports determines the state of the IP network interface associated with that VLAN. When a port-based VLAN or an IPv4 or IPv6 protocol-based VLAN comes up because one or more of its ports is up, the IP interface for that VLAN is also activated. Likewise, when a VLAN is deactivated because all of its ports are down, the corresponding IP interface is also deactivated.

VLAN MAC address

The switches have one unique MAC address for all of their VLAN interfaces. You can send an 802.2 test packet to this MAC address to verify connectivity to the switch and you can assign an IP address to the VLAN interface. When you Ping that address, ARP will resolve the IP address to this single MAC address.

In a topology where a switch has multiple VLANs and must be connected to a device having a single forwarding database, some cabling restrictions apply. For more on this topic, see **Multiple VLAN considerations** on page 27.

Port trunks

When assigning a port trunk to a VLAN, all ports in the trunk are automatically assigned to the same VLAN. A port trunk is tagged, untagged, or excluded from a VLAN the same way as individual, untrunked ports.

Port monitoring

If you designate a port on the switch for network monitoring, the port will appear in the Port VLAN Assignment screen and can be configured as a member of any VLAN. For information on how broadcast, multicast and unicast packets are tagged inside and outside of the VLAN to which the monitor port is assigned, see the section titled "VLAN-Related Problems" in the "Troubleshooting" appendix of the *ArubaOS-Switch Management and Configuration Guide* for your switch.

Jumbo packet support

Jumbo packet support is enabled per-VLAN and applies to all ports belonging to the VLAN. For more information, see the chapter titled "Port Traffic Controls" in the *ArubaOS-Switch Management and Configuration Guide* for your switch.

VLAN restrictions

- A port must be a member of at least one VLAN. In the factory default configuration, all ports are assigned to the default VLAN (DEFAULT_VLAN; VID=1).
- A port can be a member of one untagged, port-based VLAN. All other port-based VLAN assignments for that port must be tagged. The "Untagged" designation enables VLAN operation with non 802.1Q-compliant devices.

- A port can be an untagged member of one protocol-based VLAN of each protocol type. When assigning a port to multiple, protocol-based VLANs sharing of the same type, the port can be an untagged member of only one such VLAN.
- With routing enabled on the switch, the switch can route traffic between:
  - Multiple, port-based VLANs
  - A port-based VLAN and an IPv4 protocol-based VLAN
  - A port-based VLAN and an IPv6 protocol-based VLAN
  - An IPv4 protocol-based VLAN and an IPv6 protocol VLAN

  Other, routable, protocol-based VLANs must use an external router to move traffic between VLANs. With routing disabled, all routing between VLANs must be through an external router.
- Before deleting a static VLAN, t first reassign all ports in the VLAN to another VLAN. You can use the `no vlan <vid>` command to delete a static VLAN. For more information, see **Creating a new static VLAN (port-based or protocol-based) (CLI)** on page 32.
- Protocol-based VLANs, port-based VLANs and LLDP radio port VLANs cannot run concurrently with RPVST +.

Migrating Layer 3 VLANs using VLAN MAC configuration

Switches provide for maintaining Layer 3 VLAN configurations when migrating distribution routers in networks not centrally managed, by configuring the MAC address of the previous router on the VLAN interfaces of the routing switch.

VLAN MAC address reconfiguration

Switches use one unique MAC address for all VLAN interfaces. If you assign an IP address to a VLAN interface, ARP resolves the IP address to the MAC address of the routing switch for all incoming packets.

The Layer 3 VLAN MAC Configuration feature lets you reconfigure the MAC address used for VLAN interfaces, using the CLI. Packets addressed to the reconfigured Layer 3 MAC address, such as ARP and IP data packets, are received and processed by the routing switch.

Packets transmitted from the routing switch (packets originating from the router and forwarded packets) use the original Switch MAC address as the source MAC address in Ethernet headers.

ARP reply packets use the reconfigured MAC address in both the:

- ARP Sender MAC address field
- Source MAC address field in the Ethernet frame header

When reconfiguring the MAC address, you may specify a keepalive timeout to transmit heartbeat packets that advertise the new MAC address

By configuring the MAC address of the previously installed router as the MAC address of each VLAN interface on the Switch, you can swap the physical port of a router to the Switch after the switch has been properly configured in the network.

Handling incoming and outgoing VLAN Traffic

Incoming VLAN data packets and ARP requests

These are received and processed on the routing switch according to the MAC address of the previously installed router configured for each VLAN interface.

Outgoing VLAN traffic

This uses the MAC address of the switch as the source MAC address in packet headers. The MAC address configured on VLAN interfaces is not used on outbound VLAN traffic.

When the routing switch receives an ARP request for the IP address configured on a VLAN interface, the ARP reply uses the reconfigured MAC address in both the:

- ARP Sender MAC address field
- Source MAC address field in the Ethernet frame header

When proxy ARP is enabled on a VLAN interface, the ARP reply sent for an ARP request received from VLAN devices located outside the directly connected IP subnets also contains the reconfigured MAC address in both the:

- ARP Sender MAC address field
- Source MAC address field in the Ethernet frame header

To hosts in the network, VLAN traffic continues to be routed (using the reconfigured MAC address as destination address), but outbound VLAN traffic appears to be sent from another router attached to the same subnet (using the Switch MAC address as source address) attached to the same subnet. Although it appears as an asymmetric path to network hosts, the MAC address configuration feature enables Layer 3 VLAN migration. (A successful VLAN migration is achieved because the hosts do not verify that the source MAC address and the destination MAC address are the same when communicating with the routing switch.)

Sending heartbeat packets with a configured MAC Address

On the VLAN interfaces of a routing switch, the user-defined MAC address only applies to inbound traffic. As a result, any connected switches need to learn the new address that is included in the Ethernet frames of outbound VLAN traffic transmitted from the routing switch.

If a connected switch does not have the newly configured MAC address of the routing switch as a destination in its MAC address table, it floods packets to all of its ports until a return packet allows the switch to learn the correct destination address. As a result, the performance of the switch is degraded as it tries to send Ethernet packets to an unknown destination address.

To allow connected switches to learn the user-configured MAC address of a VLAN interface, the routing switch can send periodic heartbeat-like Ethernet packets. The Ethernet packets contain the configured MAC address as the source address in the packet header. IP multicast packets or Ethernet service frames are preferred because they do not interrupt the normal operation of client devices connected on the segment.

Because the aging time of destination addresses in MAC address tables varies on network devices, you must also configure a time interval to use for sending heartbeat packets.

Heartbeat packets are sent at periodic intervals with a specific Switch unicast MAC address in the destination field. This MAC address is assigned to the Switch and is not used by other non- routers. Because the heartbeat packet contains a unicast MAC address, it does not interrupt host operation. Even if you have multiple 1-65 Static Virtual LANs (VLANs) Introducing tagged VLAN technology into networks running untagged VLANs switches connected to the network, there is no impact on network performance because each switch sends heartbeat packets with its configured MAC address as the destination address.

The format of a heartbeat packet is an extended Ethernet OUI frame with an extended OUI Ethertype (88B7) and a new protocol identifier in the 5-octet protocol identifier field.

Displaying a VLAN MAC address configuration (CLI)

Syntax:

```
show ip-recv-mac-address
```

Displaying a VLAN MAC address

```
switch# show ip-recv-mac-address

VLAN L3-Mac-Address Table

VLAN                           L3-Mac-Address                 Timeout
-------------                  ------------------------       -----------
DEFAULT_VLAN                   001635-024467                  60
VLAN2                          001635-437529                  100
```

About GVRP

GVRP (GARP VLAN Registration Protocol) is an application of GARP (Generic Attribute Registration Protocol.) It enables a switch to dynamically create 802.1Q-compliant VLANs on links with other devices running GVRP and automatically create VLAN links between GVRP-aware devices. (A GVRP link can include intermediate devices that are not GVRP-aware.) This operation reduces the chance for errors in VLAN configurations by automatically providing VID (VLAN ID) consistency across the network. After the switch creates a dynamic VLAN, the CLI `static <vlan-id>` command can be used to convert it to a static VLAN. GVRP can also be used to dynamically enable port membership in static VLANs configured on a switch.

GVRP uses GVRP BPDUs (GVRP Bridge Protocol Data Units) to advertise static VLANs; this a GVRP BPDU is called an **advertisement**. On a switch, advertisements are sent outbound from ports to the devices directly connected to those ports.

GVRP operational rules

- A dynamic VLAN must be converted to a static VLAN before it can have an IP address.
- For the switches covered in this guide, GVRP can be enabled only if `max-vlans` is set to no more than 512 VLANs.
- The total number of VLANs on the switch (static and dynamic combined) cannot exceed the current Maximum VLANs setting. For example, in the factory default state, the switch supports up to 256 VLANs. Any additional VLANs advertised to the switch will not be added unless you first increase the Maximum VLANs setting. In the global config level of the CLI, use **max-vlans**.
- Converting a dynamic VLAN to a static VLAN and then executing the `write memory` command saves the VLAN in the startup-config file and makes it a permanent part of the switch's VLAN configuration.
- Within the same broadcast domain, a dynamic VLAN can pass through a device that is notGVRP-aware. This is because a half-duplex repeater, a hub, or a switch that is not GVRP-aware will flood the GVRP (multicast) advertisement packets out all ports.
- GVRP assigns dynamic VLANs as tagged VLANs. To configure the VLAN as untagged, convert it to a static VLAN.
- Rebooting a switch on which a dynamic VLAN exists deletes that VLAN. However, the dynamic VLAN reappears after the reboot if GVRP is enabled and the switch again receives advertisements for that VLAN through a port configured to add dynamic VLANs.
- By receiving advertisements from other devices running GVRP, the switch learns of static VLANs on those other devices and dynamically (automatically) creates tagged VLANs on the links to the advertising devices. Similarly, the switch advertises its static VLANs to other GVRP-aware devices, as well as the dynamic VLANs the switch has learned.
- A GVRP-enabled switch does not advertise any GVRP-learned VLANs out of the ports on which it originally learned of those VLANs.

Example of GVRP operation

In the following example, Tagged VLAN ports on switch A and switch C advertise VLANs 22 and 33 to ports on other GVRP-enabled switches that can dynamically join the VLANs.

GVRP operation

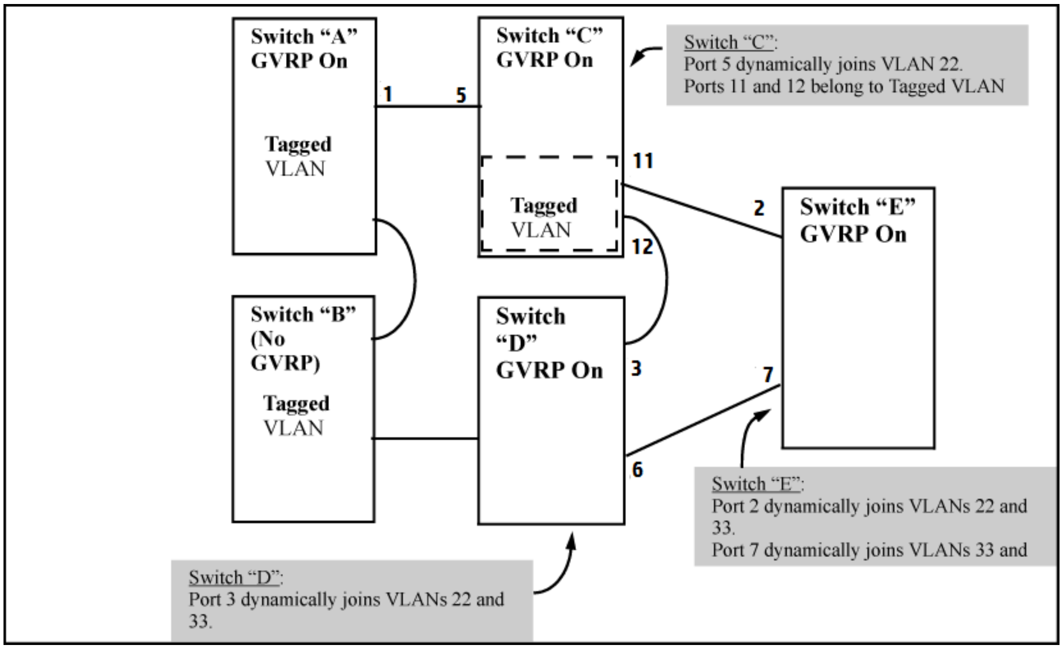

Switch "C":
Port 5 dynamically joins VLAN 22.
Ports 11 and 12 belong to Tagged VLAN

Switch "E":
Port 2 dynamically joins VLANs 22 and 33.
Port 7 dynamically joins VLANs 33 and

Switch "D":
Port 3 dynamically joins VLANs 22 and 33.

Options for a GVRP-aware port receiving advertisements

- If there is not already a static VLAN with the advertised VID on the receiving port, such a port can dynamically create the VLAN and become a member.
- If the switch already has a static VLAN assignment with the same VID as in the advertisement and the port is configured to `Auto` for that VLAN, then the port will dynamically join the VLAN and begin moving that VLAN's traffic. For more detail on `Auto`, see **Enabling a port for dynamic joins**.
- Ignore the advertisement for that VID.
- Not participate in that VLAN.

Options for a port belonging to a Tagged or Untagged static VLAN

- Send VLAN advertisements
- Receive advertisements for VLANs on other ports and dynamically join those VLANs.
- Send VLAN advertisements but ignore advertisements received from other ports.
- Avoid GVRP participation by not sending advertisements and dropping any advertisements received from other devices.

IP addressing

A dynamic VLAN does not have an IP address and moves traffic on the basis of port membership in VLANs. However, after GVRP creates a dynamic VLAN, you can convert it to a static VLAN. It is then necessary to assign ports to the VLAN in the same way that you would for a static VLAN created manually. In the static state, you can configure IP addressing on the VLAN and access it in the same way that you would any other static VLAN.

Per-port options for handling GVRP "unknown VLANs"

An "unknown VLAN" is a VLAN that the switch learns of by receiving an advertisement for that VLAN on a port that is not already a member of that VLAN. If the port is configured to learn unknown VLANs, then the VLAN is dynamically created and the port becomes a tagged member of the VLAN.

GVRP unknown VLAN settings

Suppose that in the **Example of GVRP operation**, port 1 on switch A is connected to port 5 on switch C. Because switch A has VLAN 22 statically configured, while switch C does not have this VLAN statically configured (and does not "Forbid" VLAN 22 on port 5), VLAN 22 is handled as an "Unknown VLAN" on port 5 in switch C. Conversely, if VLAN 22 was statically configured on switch C, but port 5 was not a member, port 5 would become a member when advertisements for VLAN 22 were received from switch A.

The CLI `show gvrp` command VLAN Support screen show a switch's current GVRP configuration, including the Unknown VLAN settings.

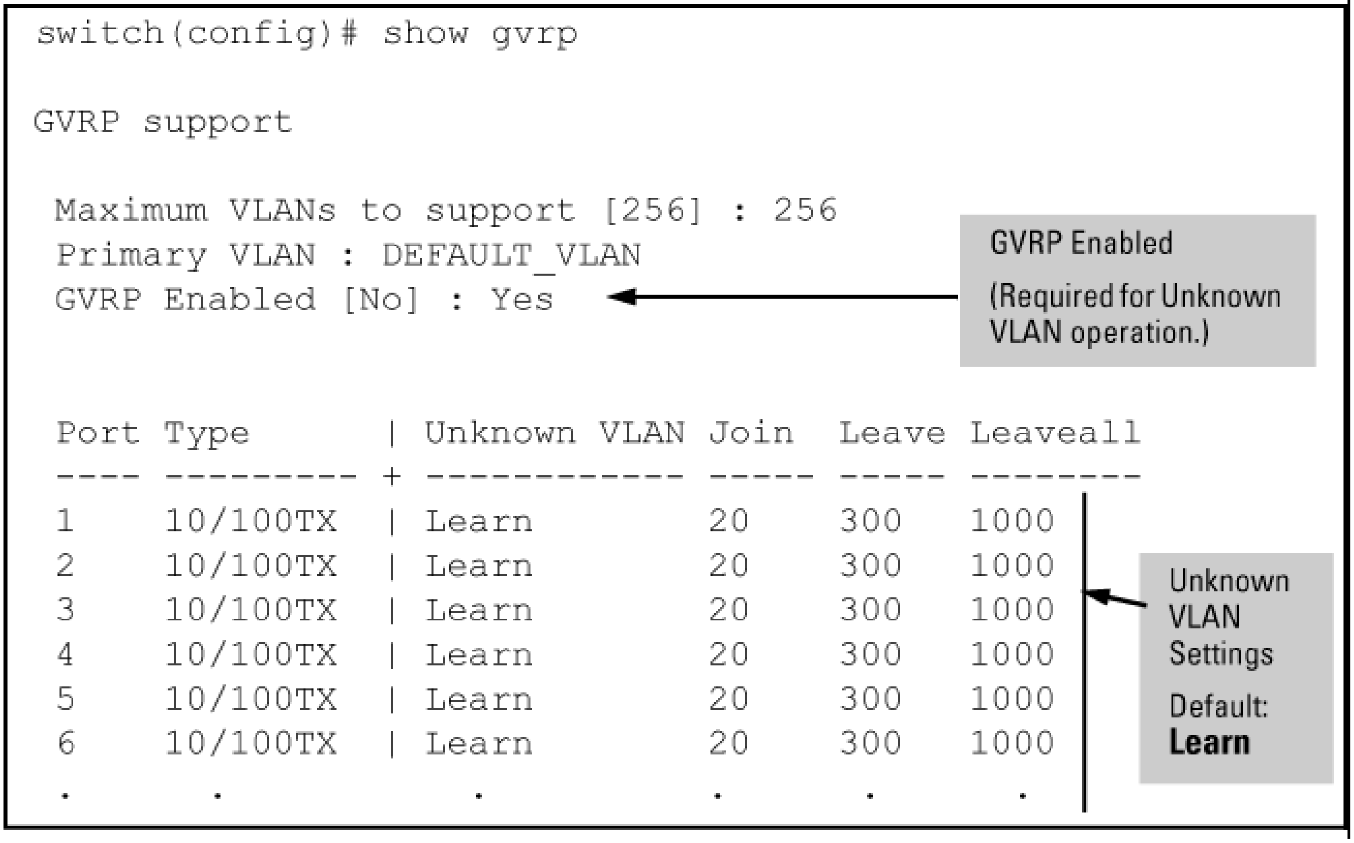

```
switch(config)# show gvrp

GVRP support

 Maximum VLANs to support [256] : 256
 Primary VLAN : DEFAULT_VLAN
 GVRP Enabled [No] : Yes


 Port Type      | Unknown VLAN Join  Leave Leaveall
 ---- --------- + ------------ ----- ----- --------
 1    10/100TX  | Learn        20    300   1000
 2    10/100TX  | Learn        20    300   1000
 3    10/100TX  | Learn        20    300   1000
 4    10/100TX  | Learn        20    300   1000
 5    10/100TX  | Learn        20    300   1000
 6    10/100TX  | Learn        20    300   1000
 .      .          .           .     .       .
```

GVRP Enabled

(Required for Unknown VLAN operation.)

Unknown VLAN Settings

Default: **Learn**

Per-port options for dynamic VLAN advertising and joining

GVRP must be enabled and VLANs must be configured to one or more switches, depending on the topology.

Initiating advertisements

As described in the preceding section, to enable dynamic joins, GVRP must be enabled and a port must be configured to Learn (the default). However, to send advertisements in your network, one or more static (Tagged, Untagged, or Auto) VLANs must be configured on one or more switches (with GVRP enabled), depending on your topology.

Enabling a port for dynamic joins

You can configure a port to dynamically join a static VLAN. The join will occur if that port subsequently receives an advertisement for the static VLAN. This is done by using the Auto and Learn options described in the table **Controlling VLAN behavior on ports with static VLANs**.

Parameters for controlling VLAN propagation behavior

You can configure an individual port to actively or passively participate in dynamic VLAN propagation or to ignore dynamic VLAN (GVRP) operation. These options are controlled by the GVRP "Unknown VLAN" and the static VLAN configuration parameters, as described in the following table.

Table 6: *Controlling VLAN behavior on ports with static VLANs*

Per-Port "Unknown VLAN" (GVRP) configuration	Static VLAN Options—Per VLAN Specified on Each Port[1]		
	Port Activity: Tagged or Untagged (Per VLAN)[2]	**Port Activity: Auto[2] (Per VLAN)**	**Port Activity: Forbid (Per VLAN)[2]**
Learn (the Default)	The port: • Belongs to specified VLAN. • Advertises specified VLAN. • Can become a member of dynamic VLANs for which it receives advertisements. • Advertises dynamic VLANs that have at least one other port (on the same switch) as a member.	The port: • Will become a member of specified VLAN if it receives advertisements for specified VLAN from another device. • Will advertise specified VLAN. • Can become a member of other, dynamic VLANs for which it receives advertisements. • Will advertise a dynamic VLAN that has at least one other port (on the same switch) as a member.	The port: • Will not become a member of the specified VLAN. • Will not advertise specified VLAN. • Can become a member of other dynamic VLANs for which it receives advertisements. • Will advertise a dynamic VLAN that has at least one other port on the same switch as a member.
Block	The port: • Belongs to the specified VLAN. • Advertises this VLAN. • Will not become a member of new dynamic VLANs for which it receives advertisements. • Will advertise dynamic VLANs that have at least one other port as a member.	The port: • Will become a member of specified VLAN if it receives advertisements for this VLAN. • Will advertise this VLAN. • Will not become a member of new dynamic VLANs for which it receives advertisements. • Will advertise dynamic VLANs that have at least one other port (on the same switch) as a member.	The port: • Will not become a member of this VLAN. • Will ignore GVRP PDUs. • Will not join any dynamic VLANs. • Will not advertise VLANs.
Disable	The port: • Is a member of the specified VLAN. • Will ignore GVRP PDUs.	The port:	The port: • Will not become a member of this VLAN. • Will ignore GVRP PDUs.

Per-Port "Unknown VLAN" (GVRP) configuration	Static VLAN Options—Per VLAN Specified on Each Port[1]		
	Port Activity: Tagged or Untagged (Per VLAN)[2]	**Port Activity: Auto[2] (Per VLAN)**	**Port Activity: Forbid (Per VLAN)[2]**
	• Will not join any advertised VLANs. • Will not advertise VLANs.	• Will not become a member of the specified VLAN. • Will ignore GVRP PDUs. • Will not join any dynamic VLANs. • Will not advertise VLANs.	• Will not join any dynamic VLANs. • Will not advertise VLANs.

[1] Each port of the switch must be a Tagged or Untagged member of at least one VLAN. Thus, any port configured for GVRP to Learn or Block will generate and forward advertisements for static VLAN(s) configured on the switch and also for dynamic VLANs the switch learns on other ports.

[2] To configure tagging, Auto, or Forbid, see **Configuring or changing static VLAN per-port settings (CLI)** on page 34.

As the preceding table indicates, when you enable GVRP, a port that has a Tagged or Untagged static VLAN has the option for both generating advertisements and dynamically joining other VLANs.

NOTE: In the table above, the Unknown VLAN parameters are configured on a per-port basis using the CLI. The Tagged, Untagged, Auto, and Forbid options are configured per static VLAN on every port.

Because dynamic VLANs operate as Tagged VLANs and because a tagged port on one device cannot communicate with an untagged port on another device, Hewlett Packard Enterprise recommends that you use Tagged VLANs for the static VLANs you will use to generate advertisements.

GVRP and VLAN access control

Enabling GVRP allows a port to advertise and join dynamic VLANs. If a port has not received an advertisement for an existing dynamic VLAN during the time-to-live (10 seconds), the port removes itself from that dynamic VLAN.

Advertisements and dynamic joins

When you enable GVRP on a switch, the default GVRP parameter settings allow all of the switch's ports to transmit and receive dynamic VLAN advertisements (GVRP advertisements) and to dynamically join VLANs.

Enabling GVRP:

- Allows a port to both advertise and join dynamic VLANs (Learn mode—the default).
- Allows a port to send VLAN advertisements, but not receive them from other devices; that is, the port cannot dynamically join a VLAN but other devices can dynamically join the VLANs it advertises (Block mode).
- Prevents a port from participating in GVRP operation (Disable mode).

Port-Leave from a dynamic VLAN

A dynamic VLAN continues to exist on a port for as long as the port receives its advertisements from another device connected to that port, or until:

- Converting the VLAN to a static VLAN
- Reconfiguring the port to `Block` or `Disable`
- Disabling GVRP
- Rebooting the switch.

The time-to-live for dynamic VLANs is 10 seconds, if a port has not received an advertisement for an existing dynamic VLAN during that time, the port removes itself from that dynamic VLAN.

Using GVRP

When GVRP is enabled on a switch, the VID for any static VLAN configured on the switch is advertised, using BPDUs (Bridge Protocol Data Units), out all ports regardless of whether a port is up or assigned to any particular VLAN. A GVRP-aware port on another device that receives the advertisements over a link can dynamically join the advertised VLAN.

A dynamic VLAN (that is, a VLAN learned through GVRP) is tagged on the port on which it was learned. Also, a GVRP-enabled port can forward an advertisement for a VLAN it learned about from other ports on the same switch (internal source), but the forwarding port will not itself join that VLAN until an advertisement for that VLAN is received through a link from another device (external source) on that specific port.

Figure 10: *Forwarding advertisements and dynamic joining*

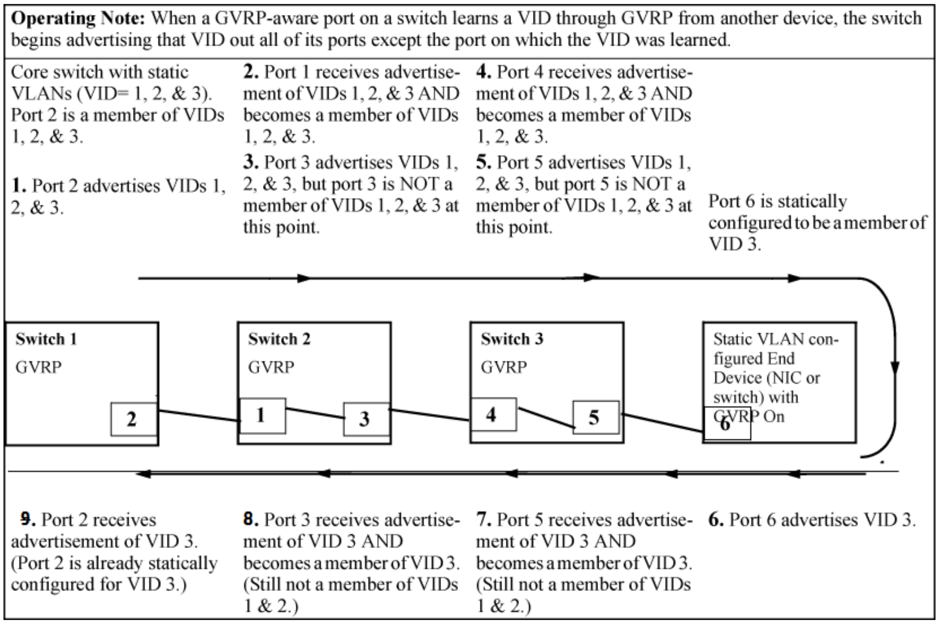

Operating Note: When a GVRP-aware port on a switch learns a VID through GVRP from another device, the switch begins advertising that VID out all of its ports except the port on which the VID was learned.

Core switch with static VLANs (VID= 1, 2, & 3). Port 2 is a member of VIDs 1, 2, & 3.

1. Port 2 advertises VIDs 1, 2, & 3.

2. Port 1 receives advertisement of VIDs 1, 2, & 3 AND becomes a member of VIDs 1, 2, & 3.

3. Port 3 advertises VIDs 1, 2, & 3, but port 3 is NOT a member of VIDs 1, 2, & 3 at this point.

4. Port 4 receives advertisement of VIDs 1, 2, & 3 AND becomes a member of VIDs 1, 2, & 3.

5. Port 5 advertises VIDs 1, 2, & 3, but port 5 is NOT a member of VIDs 1, 2, & 3 at this point.

Port 6 is statically configured to be a member of VID 3.

Switch 1
GVRP
2
Switch 2
GVRP
1
3
Switch 3
GVRP
4
5
Static VLAN configured End Device (NIC or switch) with GVRP On
6

9. Port 2 receives advertisement of VID 3. (Port 2 is already statically configured for VID 3.)

8. Port 3 receives advertisement of VID 3 AND becomes a member of VID 3. (Still not a member of VIDs 1 & 2.)

7. Port 5 receives advertisement of VID 3 AND becomes a member of VID 3. (Still not a member of VIDs 1 & 2.)

6. Port 6 advertises VID 3.

If a static VLAN is configured on at least one switch port and that port has established a link with another device, then all other ports of that switch will send advertisements for that VLAN.

NOTE:

A port can learn of a dynamic VLAN through devices that are not aware of GVRP. VLANs must be disabled in GVRP-unaware devices to allow tagged packets to pass through.

Planning for GVRP operation

To set up dynamic VLANs for a segment:

Procedure

1. Determine the VLAN topology required for each segment (broadcast domain) on the network.
2. Determine which VLANs must be static and which can be dynamically propagated.
3. Determine the devices on which static VLANs must be manually created to propagate VLANs throughout the segment.
4. Determine security boundaries and how individual ports in the segment are to handle dynamic VLAN advertisements (see **Options for handling unknown VLAN advertisements** and **Controlling VLAN behavior on ports with static VLANs**).
5. Enable GVRP on all devices to be used with dynamic VLANs and configure the appropriate "Unknown VLAN" parameter (Learn, Block, or Disable) for each port.
6. Configure static VLANs on the switches, where needed, with their per-VLAN parameters (Tagged, Untagged, Auto, and Forbid—see **Options for handling unknown VLAN advertisements** and **Controlling VLAN behavior on ports with static VLANs**) on each port.
7. Dynamic VLANs will then appear automatically, according to the chosen configuration options.
8. Convert dynamic VLANs to static VLANs, where dynamic VLANs are to become permanent.

Displaying switch current GVRP configuration (CLI)

Syntax:

```
show gvrp
```

Shows GVRP status (enabled or disabled), current maximum number of VLANs supported and the current Primary VLAN.

Displaying GVRP status with GVRP disabled

```
switch(config)# show gvrp

GVRP support

 Maximum VLANs to support [256] : 256
 Primary VLAN : DEFAULT_VLAN
 GVRP Enabled [No] : No
```

Displaying GVRP status with GVRP enabled

This example shows the output for the `show gvrp` command with GVRP enabled. It includes non-default settings for the Unknown VLAN field for some ports (see Port number 3, 4, 5 below).

```
switch(config)# show gvrp

 GVRP support

  Maximum VLANs to support [256] : 256
  Primary VLAN : DEFAULT_VLAN
  GVRP Enabled [No] : Yes
```

```
Port Type      | Unknown VLAN Join  Leave Leaveall
---- --------- + ------------ ----- ----- --------
1    10/100TX  | Learn        20    300   1000
2    10/100TX  | Learn        20    300   1000
3    10/100TX  | Block        20    300   1000
4    10/100TX  | Disable      20    300   1000
5    10/100TX  | Disable      20    300   1000
6    10/100TX  | Learn        20    300   1000
7    10/100TX  | Learn        20    300   1000
```

Displaying switch current GVRP configuration (CLI)

Syntax:

```
show gvrp
```

Shows GVRP status (enabled or disabled), current maximum number of VLANs supported and the current Primary VLAN.

Displaying GVRP status with GVRP disabled

```
switch(config)# show gvrp

GVRP support

 Maximum VLANs to support [256] : 256
 Primary VLAN : DEFAULT_VLAN
 GVRP Enabled [No] : No
```

Displaying GVRP status with GVRP enabled

This example shows the output for the `show gvrp` command with GVRP enabled. It includes non-default settings for the Unknown VLAN field for some ports (see Port number 3, 4, 5 below).

```
switch(config)# show gvrp

 GVRP support

  Maximum VLANs to support [256] : 256
  Primary VLAN : DEFAULT_VLAN
  GVRP Enabled [No] : Yes


Port Type      | Unknown VLAN Join  Leave Leaveall
---- --------- + ------------ ----- ----- --------
1    10/100TX  | Learn        20    300   1000
2    10/100TX  | Learn        20    300   1000
3    10/100TX  | Block        20    300   1000
4    10/100TX  | Disable      20    300   1000
5    10/100TX  | Disable      20    300   1000
6    10/100TX  | Learn        20    300   1000
7    10/100TX  | Learn        20    300   1000
```

Enabling and disabling GVRP on the switch (CLI)

Syntax:

```
gvrp
```

Enables GVRP on the switch.

```
no gvrp
```

Disables GVRP on the switch.

NOTE:

GVRP can be enabled only if `max-vlans` is set to no more than 256 VLANs. While GVRP is enabled on the switch, you cannot apply any ACLs to VLANs configured on the same switch. A GVRP link can include intermediate devices that are not GVRP-aware. To understand and use GVRP, you need a working knowledge of 802.1Q VLAN tagging. See **802.1Q VLAN tagging** on page 19.

GVRP assigns dynamic VLANs as Tagged VLANs. To configure the VLAN as Untagged, you must first convert it to a static VLAN.

A VLAN enabled for jumbo traffic cannot be used to create a dynamic VLAN. A port belonging to a statically configured, jumbo-enabled VLAN cannot join a dynamic VLAN.

Controlling how individual ports handle advertisements for new VLANs (CLI)

When GVRP is enabled on the switch, use the `unknown-vlans` command to change the Unknown VLAN field for one or more ports.

Syntax:

```
interface <port-list> unknown-vlans [learn | block | disable]
```

Changes the Unknown VLAN field to control how one or more ports handle advertisements. Use at either the Manager or interface context level for a port.

Changing the Unknown VLANs field

In the following example, the first command changes the configuration to Block, the second command displays the new configuration:

```
switch(config)# interface 1-2 unknown-vlans block

Switch(config)# show gvrp
 GVRP support
  Maximum VLANs to support [256] : 256
  Primary VLAN : DEFAULT_VLAN
  GVRP Enabled [No] : Yes

  Port Type      | Unknown VLAN Join  Leave Leaveall
  ---- --------- + ------------ ----- ----- --------
  1    10/100TX  | Block        20    300   1000
  2    10/100TX  | Block        20    300   1000
  3    10/100TX  | Learn        20    300   1000
  4    10/100TX  | Learn        20    300   1000
```

When you enable GVRP on a switch, you have the per-port join-request options listed in the following table:

Table 7: *Options for handling unknown VLAN advertisements*

Unknown VLAN Mode	Operation
`Learn` (the Default)	Enables the port to become a member of any unknown VLAN for which it receives an advertisement. Allows the port to advertise other VLANs that have at least one other port on the same switch as a member.
`Block`	Prevents the port from joining any new dynamic VLANs for which it receives an advertisement.Allows the port to advertise other VLANs that have at least one other port as a member.
`Disable`	Causes the port to ignore and drop all GVRP advertisements it receives and prevents the port from sending any GVRP advertisements.

Listing static and dynamic VLANs on a GVRP-enabled switch (CLI)

Syntax:

```
show vlans
```

Lists all VLANs present in the switch.

Using the `show vlans` command

In the following illustration, switch B has one static VLAN (the default VLAN), with GVRP enabled and port 1 configured to `Learn` for Unknown VLANs. Switch A has GVRP enabled and has three static VLANs: the default VLAN, VLAN-222 and VLAN-333. In this scenario, switch B will dynamically join VLAN-222 and VLAN-333:

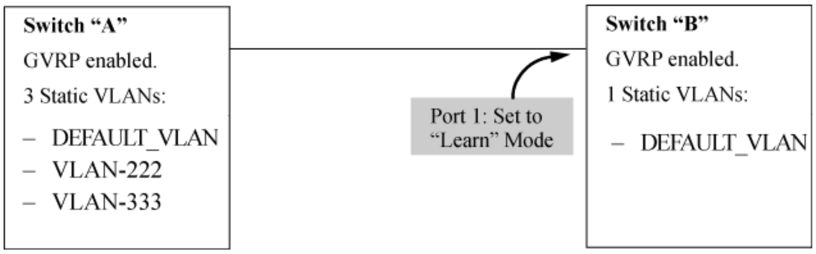

The `show vlans` command lists the dynamic (and static) VLANs in switch B after it has learned and joined VLAN-222 and VLAN-333.

```
Switch-B> show vlans

 Status and Counters - VLAN Information

   VLAN support : Yes
   Maximum VLANs to support : 16
   Primary VLAN : DEFAULT_VLAN

   VLAN ID        NAME          Status
   -------------- ------------- ------
   1              DEFAULT_VLAN  Static
   222            GVRP_222       Dynamic
   333            GVRP_333       Dynamic
```

Converting a Dynamic VLAN to a Static VLAN (CLI)

If a port on the switch has joined a dynamic VLAN, you can use the following command to convert that dynamic VLAN to a static VLAN:

```
static-vlan <dynamic-vlan-id>
```

Converting a dynamic VLAN 333 to a static VLAN

When converting a dynamic VLAN to a static VLAN as shown here, all ports on the switch are assigned to the VLAN in Auto mode.

```
switch(config)# static-vlan 333
```

Multiple VLAN Registration Protocol overview

Multiple VLAN Registration Protocol (MVRP) is a registration protocol defined by IEEE, which propagates VLAN information dynamically across devices. It also enables devices to learn and automatically synchronize VLAN configuration information, thereby reducing the configuration workload.

It is an enhanced version of GVRP and improves declaration efficiency. It allows a participant (port) to make or withdraw declaration of attributes (VLANs). These declarations (or withdraws) are resulted in registration (or removal of registrations) with other switches in the network.

Salient features

- Complaint as per IEEE 802.1Q-2011(Clause 11.2).
- Supports conversion of dynamic VLAN to static VLAN.
- Supports propagation of radius assigned dynamic VLANs.
- Supports immediate registration and propagation of VLAN attributes during spanning tree topology changes.
- Supports registrar’s administrative control values such as normal, fixed, and forbid.
- Supports MVRP objects on the following standard MIBs:
  - IEEE8021-Q-BRIDGE-MIB (version 200810150000Z)
  - IEEE8021-BRIDGE-MIB (version 200810150000Z)

NOTE: Supports other MVRP objects with the help of proprietary MIB, HPE-ICF-MVRP-MIB `(hpicfMvrp.mib).`

- Supports on both physical and LAG ports, which include the manual (trunk), static lacp, and dynamic lacp trunks.
- Supports High Availability hitless.
- Supports configuring MVRP using CLI and SNMP commands.
- Supports configurable timers — Join, Leave, Leave-All, and Periodic.
- Supports fast logging for important MVRP events and error conditions.
- Supports debug logging for all MVRP enabled ports.
- MVRP can be used to manage VLANs on dynamic trunk.

MVRP operating notes

MVRP is an enhanced version of Generic Attribute Registration Protocol (GARP). It is a generic registration framework defined by the IEEE 802.1ak amendment to the IEEE 802.1Q standard. As GVRP, the same rules for dynamic propagation and registration of VLANs is also applicable for MVRP on Aruba switches.

- A dynamic VLAN must be converted to a static VLAN before it can have an IP address.
- On the switches covered in this guide, MVRP can be enabled only if `max-vlans` is not more than 512 VLANs.
- The total number of VLANs on the switch (static and dynamic combined) cannot exceed the current maximum VLANs setting. For example, in the factory default state, the switch supports up to 256 VLANs. Any additional VLANs advertised to the switch are not added unless you increase the maximum VLANs setting.
- Converting a dynamic VLAN to a static VLAN and then executing the `write memory` command saves the VLAN in the `startup-config` file and makes it a permanent part of the switch's VLAN configuration.
- When you enable MVRP globally, it is enabled by default on dynamic trunks. Based on your requirement, you can disable MVRP on dynamic trunks. You cannot modify any other MVRP port parameters.
- Within the same broadcast domain, a dynamic VLAN can pass through a device that is not MVRP-aware. This is because a half-duplex repeater or a switch that is not MVRP-aware floods the MVRP (multicast) advertisement packets out of all ports.
- Rebooting a switch on which a dynamic VLAN exists deletes the VLAN. However, the dynamic VLAN reappears after the reboot, if MVRP is enabled. The switch again receives advertisement for the particular VLAN through a port configured to add dynamic VLANs.
- By receiving advertisements from other devices running MVRP, the switch learns of static VLANs on those devices and dynamically (automatically) creates tagged VLANs on the links to the advertising devices. Similarly, the switch advertises its static VLANs and the dynamic VLANs to other MVRP-aware devices, which the switch has learnt.
- An MVRP enabled switch does not advertise any MVRP learned VLANs out of the ports (on which it originally learned of those VLANs), until it is dynamically learnt on at least two ports.
- While MVRP is enabled on the switch, you cannot apply any ACLs to VLANs configured on the same switch.

Listing static and dynamic VLANs on an MVRP-enabled switch

Syntax

`show vlan`

Description

Displays both static and dynamic VLANs in the switch.

Example output

```
switch(config)# show vlan

Status and Counters - VLAN Information

  Maximum VLANs to support : 256
  Primary VLAN : DEFAULT_VLAN
  Management VLAN :

  VLAN ID Name                             | Status     Voice Jumbo
  ------- -------------------------------- + ---------- ----- -----
  1       DEFAULT_VLAN                     | Port-based No    No
  40      MVRP_40                          | Dynamic
```

Converting a dynamic VLAN to a static VLAN

Syntax

```
static-vlan <dynamic-vlan-id>
```

Description

If a port on the switch has joined a dynamic VLAN, use the command to convert dynamic VLAN to static VLANs in the switch.

Example output

```
switch(config)# static-vlan 40
switch(config)# show vlan

Status and Counters - VLAN Information

  Maximum VLANs to support : 256
  Primary VLAN : DEFAULT_VLAN
  Management VLAN :

  VLAN ID Name                             | Status     Voice Jumbo
  ------- -------------------------------- + ---------- ----- -----
  1       DEFAULT_VLAN                     | Port-based No    No
  40      VLAN40                           | Port-based No    No
```

Viewing the current MVRP configuration on a switch

show mvrp

Syntax

```
show mvrp [config|state|statistics]
```

Description

Displays the MVRP settings and status.

Example output

```
switch# show mvrp
 config                Show the MVRP configuration for all ports.
 state                 Show the MVRP state.
 statistics            Show MVRP statistics.
```

show mvrp config

Syntax

```
show mvrp config
```

Description

Displays the MVRP configuration for all ports.

Example output

```
switch# show mvrp config

 Configuration and Status - MVRP

 Global MVRP status : Disabled

 Port     Status   Periodic Registration Join  Leave    LeaveAll Periodic
```

```
                   Timer    Type          Time  Timer    Timer    Timer
 -------  --------- -------- ------------ ----  -------- -------- --------
 1        Disabled Enabled  Normal       20    300      1000     100
 2        Disabled Enabled  Normal       20    300      1000     100
 3        Disabled Enabled  Normal       20    300      1000     100
```

show mvrp state

Syntax

`show mvrp state <VLAN-ID> [<PORT-NUM>]`

Description

Displays the MVRP state.

Parameters

<VLAN-ID>

Specify the MVRP state for VLAN ID.

<PORT-NUM>

Specify the port number to display the MVRP state.

Example output

```
switch(config)# show mvrp state
 VLAN-ID                Enter a VLAN identifier or the VLAN name if configured.
switch(config)# show mvrp state 1
 [ethernet] PORT-NUM
switch(config)# show mvrp state 1

 Configuration and Status - MVRP state for VLAN 1

 Port     VLAN  Registrar Applicant Forbid
                State     State     Mode
 -------- ----- --------- --------- ---------
 1        1     MT        QA        No
```

show mvrp statistics

Syntax

`show MVRP statistics [<PORT-LIST>]`

Description

Displays the MVRP statistics.

Parameter

`PORT-LIST`

Displays the MVRP statistics at the specified port.

Example output

```
switch(config)# show mvrp statistics

Status and Counters - MVRP

  MVRP statistics for port  : A1
  ----------------------------
  Failed registration    : 0
  Last PDU origin        : 40a8f0-9e11ff
  Total PDU Transmitted  : 53
  Total PDU Received     : 72
  Frames Discarded       : 0

  Message type  Transmitted     Received
-------------- ------------    ------------
  New           0               0
  Empty         16466           258
  In            4               0
  Join Empty    0               72
  Join In       53              55
  Leave         0               0
  Leaveall      4               2
```

clear mvrp statistics

Syntax

`clear mvrp statistics [<PORT-LIST>]`

Description

Clears the statistics for MVRP on a port or all ports.

Parameters

`PORT-LIST`

Specify a port number or list of ports or all ports.

Example output

```
switch# clear mvrp statistics
 [ethernet] PORT-LIST  Enter a port number, a list of ports or 'all' for all ports.
switch# clear mvrp statistics all
```

debug mvrp

Syntax

`debug mvrp {all | event| packet | state-machine | timer} [<PORT-LIST>]`

Description

Enables debug messages.

Parameters

`all`

Display all MVRP debug messages.

`event`

Display all MVRP event messages.

`packet`

Display all MVRP packet messages.

`state-machine`

Display all MVRP state-machine messages.

`timer`

Display all MVRP timer messages.

`PORT-LIST`

Display all MVRP debug messages for a port.

Example output

```
switch(config)# debug mvrp all
switch(config)# show debug

Debug Logging

  Source IP Selection: Outgoing Interface
  Origin identifier: Outgoing Interface IP
  Destination: None

  Enabled debug types:
   mvrp event include port A1-A24,F1-F24
   mvrp packet include port A1-A24,F1-F24
   mvrp state-machine include port A1-A24,F1-F24
   mvrp timer include port A1-A24,F1-F24
```

Configuring MVRP

Enabling MVRP globally

MVRP must be enabled globally to allow the device to participate in the protocol.

Syntax

`mvrp {enable | disable}`

`no mvrp`

Description

Enables MVRP globally on a switch. MVRP must be enabled globally and at least on one interface. The `no` form of the command disables MVRP.

Parameters

`enable`

Enable MVRP.

`disable`

Disable MVRP.

Example output

```
switch# show mvrp config

 Configuration and Status - MVRP

 Global MVRP status : Enabled

 Port     Status   Periodic Registration Join  Leave    LeaveAll Periodic
                   Timer    Type         Time  Timer    Timer    Timer
 -------  -------- -------- ------------ ----  -------- -------- --------
 1        Disabled Enabled  Normal       20    300      1000     100
 2        Disabled Enabled  Normal       20    300      1000     100
```

Enabling MVRP on an interface

By default, MVRP is disabled on all interfaces.

Syntax

`mvrp {enable | disable}`

`no mvrp`

Description

Enables MVRP on an interface. MVRP must be enabled globally and at least on one interface.

Use `no mvrp` to disable MVRP.

Parameters

`enable`

Enable `mvrp`

`disable`

Disable `mvrp`

Example output

```
switch(config)# mvrp
disable               Disable MVRP.
enable                Enable MVRP.
switch(config)# mvrp enable
switch(config)# interface 1
switch(eth-1)# mvrp enable
switch(eth-1)# show mvrp config

 Configuration and Status - MVRP

 Global MVRP status : Enabled

 Port     Status   Periodic Registration Join  Leave    LeaveAll Periodic
                   Timer    Type         Time  Timer    Timer    Timer
 -------  -------- -------- ------------ ----  -------- -------- --------
 1        Enabled  Enabled  Normal       20    300      1000     100
 2        Disabled Enabled  Normal       20    300      1000     100
```

MVRP timers

MVRP supports four types of timers:

- Join Timer
- Leave Timer
- LeaveAll Timer
- Periodic Timer

Join Timer

The Join Timer controls the transmission of Join messages. To avoid a PDU storm, an MVRP participant waits for a duration of the Join Timer after sending a join message, and ensures that all participants transmit at different times. This is a per port timer and is applicable to all applicants for the port.

mvrp join-timer

Syntax

`mvrp join-timer <centiseconds>`

`no mvrp join-timer`

Description

Sets the Join Timer for the port. You can use the timer to space MVRP join messages. To ensure that join messages are transmitted to other participants, an MVRP participant waits for a specified time before sending a join message. The Join Timer must be less than half of the Leave Timer. The default value is 20 centiseconds.

Use `no mvrp join-timer` to set the interval to the default value.

Parameters

`centiseconds`

Set the Join Timer for the port.

Usage

`mvrp join-timer <20-100>`

The MVRP Join Timer ranges from 20 –100 in centiseconds.

Example output

```
switch(eth-1)# mvrp join-timer
<20-100>              Set the join timer for the port.
switch# mvrp join-timer 40
switch# show mvrp config

 Configuration and Status - MVRP

 Global MVRP status : Enabled

 Port     Status   Periodic Registration Join  Leave    LeaveAll Periodic
                   Timer    Type         Time  Timer    Timer    Timer
 -------  -------- -------- ------------ ----  -------- -------- --------
 1        Enabled  Enabled  Normal       40    300      1000     100
 2        Disabled Enabled  Normal       20    300      1000     100
 3        Disabled Enabled  Normal       20    300      1000     100
```

Leave Timer

The Leave Timer controls the time duration for which the Registrar state machine waits in the LV state before changing to the MT state. The Leave Timer is started only when a leave message is received by the applicant state. The attribute is deregistered, if there are requests to join before the expiry of the Leave Timer. This is a per port timer and is applicable to all registrars for the port.

mvrp leave-timer

Syntax

`mvrp leave-timer <centiseconds>`

`no mvrp leave-timer`

Description

The Leave Timer must be at least twice the Join Timer and must be less than the LeaveAll Timer. The default value is 300 centiseconds.

Use `no mvrp leave-timer` to set the interval to the default value.

Parameter

`centiseconds`

Set the Leave Timer for the port.

Usage

`mvrp leave-timer <40-1000000>`

The MVRP Leave Timer ranges from 40 –1000000 in centiseconds.

Example output

```
switch(eth-1)# mvrp leave-timer
<40-1000000>          Set the leave timer for the port.
switch(eth-1)# mvrp leave-timer 500
switch(eth-1)# show mvrp config

 Configuration and Status - MVRP

 Global MVRP status : Enabled

 Port     Status   Periodic Registration Join  Leave    LeaveAll Periodic
                   Timer    Type         Time  Timer    Timer    Timer
 -------  -------- -------- ------------ ----  -------- -------- --------
 1        Enabled  Enabled  Normal       40    500      1000     100
 2        Disabled Enabled  Normal       20    300      1000     100
 3        Disabled Enabled  Normal       20    300      1000     100
```

LeaveAll Timer

The LeaveAll Timer controls the frequency with which the LeaveAll state machine generates LeaveAll PDUs. When a LeaveAll Timer expires, the MVRP sends out LeaveAll messages and restarts the LeaveAll Timer. The LeaveAll Timer is set to a random value `T` which ranges from `LeaveAllTime < T < 1.5*LeaveAllTime`, where LeaveAll time is the configured LeaveAll time. The default value is 1000 centiseconds. This is a per port timer.

mvrp leaveall-timer

Syntax

`mvrp leaveall-timer <centiseconds>`

`no mvrp leaveall-timer`

Description

The LeaveAll Timer is the time duration between sending LeaveAll messages. The LeaveAll Timer must be greater than the Leave Timer.

Use `no mvrp leaveall-timer` to set the interval to the default value.

Parameter

`centiseconds`

Set the LeaveAll Timer for the port.

Usage

`mvrp leaveall-timer <500-1000000>`

The MVRP LeaveAll Timer ranges from 500 –1000000 in centiseconds.

Example output

```
switch# mvrp leaveall-timer
<500-1000000> Set the leaveall timer for the port.
switch# mvrp leaveall-timer 700
switch(eth-1)# show mvrp config

 Configuration and Status - MVRP

 Global MVRP status : Enabled
 Port     Status   Periodic Registration Join  Leave    LeaveAll Periodic
                   Timer    Type         Time  Timer    Timer    Timer
 -------  -------- -------- ------------ ----  -------- -------- --------
 1        Enabled  Enabled  Normal       40    500      700      100
 2        Disabled Enabled  Normal       20    300      1000     100
 3        Disabled Enabled  Normal       20    300      1000     100
```

Periodic Timer

The Periodic Timer controls the frequency with which the periodic transmission state machine generates periodic events. This is a per port timer. On start, the Periodic Timer is set to one second. You can enable or disable the Periodic Timer. By default, it is enabled. The default value is 100 centiseconds.

mvrp periodic timer

Syntax

`mvrp periodic-timer <centiseconds>`

`no mvrp periodic-timer`

Description

Set the Periodic Timer transmission interval for the port.

Use `no mvrp periodic-timer` to set the interval to the default value.

Parameters

`centiseconds`

Set the Periodic Timer transmission interval for the port.

Usage

`mvrp periodic-timer <100-1000000>`

The MVRP Periodic Timer ranges from 100 –1000000 in centiseconds.

Example output

```
switch(eth-1)# mvrp periodic-timer
<100-1000000>         Set the periodic timer transmission interval for the port.
switch(eth-1)# mvrp periodic-timer 300
switch(eth-1)# show mvrp config

 Configuration and Status - MVRP

 Global MVRP status : Enabled

 Port     Status   Periodic Registration Join  Leave    LeaveAll Periodic
                   Timer    Type         Time  Timer    Timer    Timer
 -------  -------- -------- ------------ ----  -------- -------- --------
 1        Enabled  Enabled  Normal       40    500      700      300
 2        Disabled Enabled  Normal       20    300      1000     100
 3        Disabled Enabled  Normal       20    300      1000     100
```

mvrp periodic-timer-enable

Syntax

`mvrp periodic-timer-enable`

`no mvrp periodic-timer-enable`

Description

Enable Periodic Timer transmission for the port. By default, it is enabled.

Use `no mvrp periodic-timer-enable` to disable the Periodic Timer on an interface.

MVRP registration modes

MVRP supports three registration modes:

- **Normal**

  In this mode, a port can register and deregister dynamic VLANs. By default, the registrar mode is normal.

- **Fixed**

  In this mode, a port cannot register or deregister dynamic VLANs. However, if a static VLAN exists in the system, the port changes to registered state on receipt of join message.

- **Forbidden**

  In this mode, a port does not register dynamic VLANs, ignores all MRP messages, and remains in MT state (unregistered).

mvrp registration

Syntax

`mvrp registration {normal |fixed}`

Description

Configures the port response to MRP messages.

Parameters

`normal`

Port response is normal for the incoming MRP messages.

`fixed`

Ignores the MRP messages and remains registered.

Example output

```
switch# mvrp registration
 fixed                 The port ignores all MRP messages and remains registered.
 normal                The port responds normally to incoming MRP messages.

switch(config)# interface A1 mvrp registration fixed
switch(config)# show mvrp config

Configuration and Status - MVRP

Global MVRP status : Enabled

Port     Status   Periodic Registration Join  Leave    LeaveAll Periodic
                   Timer    Type         Time  Timer    Timer    Timer
-------  -------- -------- ------------ ----  -------- -------- --------
A1       Enabled  Enabled  Fixed        20    300      1000     100
A2       Disabled Enabled  Normal       20    300      1000     100
A3       Disabled Enabled  Normal       20    300      1000     100
```

show tech mvrp

Syntax

`show tech mvrp`

Description

Displays statistics of all the MVRP enabled ports.

Example output

```
switch# show tech mvrp

show mvrp statistics

Status and Counters - MVRP

  MVRP statistics for port  : A1
  ----------------------------
  Failed registration    : 0
  Last PDU origin        : 40a8f0-9e11ff
  Total PDU Transmitted  : 620
  Total PDU Received     : 755
  Frames Discarded       : 0



  Message type  Transmitted     Received
-------------- ------------    ------------
  New           0               0
  Empty         117370          2506
  In            17              0
  Join Empty    1               519
  Join In       658             697
  Leave         0               0
  Leaveall      28              37

mvrpDumpGlobalData

MVRP global enabled status    : enabled
MVRP enabled ports            : A1
Total MVRP enabled ports      : 1
Dyn trunk auto disable count  : 0
Total Static VLANs in system  : 1
Total Dynamic VLANs in system : 1
Max VLANs supported           : 512

Display VLAN_GROUP to VLANs Mapping:

Group ID      Mapped VLANs
----------    ----------------
       0      1-4094

Display timer Ports:

Group ID     Timer Value
----------   -------------

Display Blocked Ports:

Group ID     Blocked Ports
----------   ---------------

mvrppconfig

Mvrp Port state info:

Port   MvrpState   LinkState   Registrar   Value
-----   ---------   ---------   ---------   -----
```

```
   A1      Enable          Up        Normal   0X05
   A2     Disable          Up        Normal   0X04
   A3     Disable        Down        Normal   0000
   A4     Disable        Down        Normal   0000
   A5     Disable        Down        Normal   0000
   A6     Disable        Down        Normal   0000
   A7     Disable        Down        Normal   0000
   A8     Disable        Down        Normal   0000
   A9     Disable        Down        Normal   0000
  A10     Disable        Down        Normal   0000
  A11     Disable        Down        Normal   0000
  A12     Disable        Down        Normal   0000
  A13     Disable        Down        Normal   0000
  A14     Disable        Down        Normal   0000
  A15     Disable        Down        Normal   0000
  A16     Disable        Down        Normal   0000
  A17     Disable        Down        Normal   0000
  A18     Disable        Down        Normal   0000
  A19     Disable        Down        Normal   0000
  A20     Disable        Down        Normal   0000
  A21     Disable        Down        Normal   0000
  A22     Disable        Down        Normal   0000
  A23     Disable        Down        Normal   0000
  A24     Disable        Down        Normal   0000
   F1     Disable        Down        Normal   0000
   F2     Disable        Down        Normal   0000
   F3     Disable        Down        Normal   0000
   F4     Disable        Down        Normal   0000
   F5     Disable        Down        Normal   0000
   F6     Disable        Down        Normal   0000
   F7     Disable        Down        Normal   0000
   F8     Disable        Down        Normal   0000
   F9     Disable        Down        Normal   0000
  F10     Disable        Down        Normal   0000
  F11     Disable        Down        Normal   0000
  F12     Disable        Down        Normal   0000
  F13     Disable        Down        Normal   0000
  F14     Disable        Down        Normal   0000
  F15     Disable        Down        Normal   0000
  F16     Disable        Down        Normal   0000
  F17     Disable        Down        Normal   0000
  F18     Disable        Down        Normal   0000
  F19     Disable        Down        Normal   0000
  F20     Disable        Down        Normal   0000
  F21     Disable          Up        Normal   0X04
  F22     Disable          Up        Normal   0X04
  F23     Disable        Down        Normal   0000
  F24     Disable        Down        Normal   0000

Mvrp Port timer values:

Port  join  leave  leaveall  periodic  periodic-enabled
-----  ----  -----  --------  --------  ----------------
   A1     20    300      1000       100          enabled
   A2     20    300      1000       100          enabled
   A3     20    300      1000       100          enabled
   A4     20    300      1000       100          enabled
   A5     20    300      1000       100          enabled
   A6     20    300      1000       100          enabled
   A7     20    300      1000       100          enabled
   A8     20    300      1000       100          enabled
   A9     20    300      1000       100          enabled
  A10     20    300      1000       100          enabled
  A11     20    300      1000       100          enabled
  A12    20    300      1000       100          enabled
  A13    20    300      1000       100          enabled
  A14    20    300      1000       100          enabled
  A15    20    300      1000       100          enabled
  A16    20    300      1000       100          enabled
  A17    20    300      1000       100          enabled
  A18    20    300      1000       100          enabled
  A19    20    300      1000       100          enabled
  A20    20    300      1000       100          enabled
  A21    20    300      1000       100          enabled
  A22    20    300      1000       100          enabled
  A23    20    300      1000       100          enabled
  A24    20    300      1000       100          enabled
   F1    20    300      1000       100          enabled
   F2    20    300      1000       100          enabled
   F3    20    300      1000       100          enabled
   F4    20    300      1000       100          enabled
   F5    20    300      1000       100          enabled
   F6    20    300      1000       100          enabled
   F7    20    300      1000       100          enabled
   F8    20    300      1000       100          enabled
   F9    20    300      1000       100          enabled
  F10    20    300      1000       100          enabled
  F11    20    300      1000       100          enabled
  F12    20    300      1000       100          enabled
  F13    20    300      1000       100          enabled
  F14    20    300      1000       100          enabled
  F15    20    300      1000       100          enabled
  F16    20    300      1000       100          enabled
  F17    20    300      1000       100          enabled
  F18    20    300      1000       100          enabled
  F19    20    300      1000       100          enabled
  F20    20    300      1000       100          enabled
  F21    20    300      1000       100          enabled
  F22    20    300      1000       100          enabled
  F23    20    300      1000       100          enabled
  F24    20    300      1000       100          enabled

mvrpmapringShow

Mvrp list info:
-------------------------
Port A1    : connected

Mvrp Map Count Info:

Vlan   Vid    Reg-Count
-----  -----  ----------
1      1      1
2      40     1

=== The command has completed successfully. ===
```

MVRP limitations

- MVRP and GVRP are mutually exclusive, and cannot coexist.
- MVRP and Smartlink are mutually exclusive. Smartlinks can be enabled on ports, which are not MVRP enabled and vice versa.
- MVRP and PVST are mutually exclusive. When MVRP is globally enabled, spanning tree mode cannot be set as PVST and vice versa.

- MVRP can be enabled on a provider bridge environment, but does not support SVLAN ports in mixed mode configuration.
- MVRP can be used to manage VLANs on dynamic trunk.
- Enable `aaa port-access gvrp-vlans` to support RADIUS-assigned VLANs. When you enable `aaa port-access gvrp-vlans`, dynamic VLANs created by MVRP or GVRP can be used for radius port assignment.
- An OpenFlow member VLAN cannot be a dynamic VLAN. As a result, a dynamic VLAN must be converted to static to be handled by the OpenFlow controller.
- For security purposes, MVRP is disabled by default. MVRP packets are blocked on MVRP disabled ports, but can be enabled on ports which are security enabled.
- MVRP and private VLAN cannot coexist.
- DIPLDv6 cannot be configured on MVRP enabled ports.
- MVRP support is limited to 512 VLANs and 24 logical ports due to CPU and memory resource availability.

Table 8: *MVRP supported ports*

Platforms	Maximum MVRP ports supported
`Aruba 2530`	24

Table 9: *MVRP supported VLANs*

Platforms	Maximum VLANs	Maximum MSTP instance	Maximum ports
`Aruba 2530`	512	16	24

MVRP statistics

The MVRP statistics generated using `show mvrp statistics`, records any registration failures, tracks MAC addresses to derive statistics.

- **Registration failure**

  Maintains the count of registration requests received but failed due to MVRP limitation.
- **Peer tracking**

  Records the MAC address of the MVRP PDU that has caused the recent state change for the registrar machine. A maximum of one MAC address per port of the originator switch is stored.
- **PDU event statistics**

  Collects the data on numbers of events (join, leave, and so on) transmitted and received.

For more information, see **show mvrp statistics**.

Operation and features

NOTE: Multicast filtering is not supported on switch models J9779A, J9780A, J9782A and J9783A.

In a network where IP multicast traffic is transmitted for multimedia applications, you can use a switch to reduce unnecessary per-port bandwidth usage by configuring IGMP (Internet Group Management Protocol) controls. In the factory default state (IGMP disabled), the switch floods all IP multicast traffic it receives on a given VLAN through all ports on that VLAN, except the port on which it received the traffic. This can cause significant and unnecessary bandwidth use in networks employing IP multicast traffic. With IGMP, ports can detect IGMP queries, report packets and manage IP switch multicast traffic.

IGMP is useful in multimedia applications such as LAN TV, desktop conferencing and collaborative computing that have multipoint communication (communication from one-to-many or many-to-many hosts). In such multipoint applications, IGMP is configured on the hosts and multicast traffic is generated by one or more servers (inside or outside the local network). Switches in the network that support IGMP can then be configured to direct the multicast traffic to only the ports where needed. If multiple VLANs are configured, you can configure IGMP by VLAN.

Enabling IGMP allows detection of IGMP queries and report packets to manage IP multicast traffic through the switch. If no other querier is detected, the switch then also functions as the querier. To disable the querier feature, use the IGMP configuration MIB (see "Configuring the querier function" in **CLI: Configuring and displaying IGMP**).

NOTE: IGMP configuration on the switch operates at the VLAN context level. If you are not using VLANs, then configure IGMP in VLAN 1 (the default VLAN) context.

IGMP devices

- **IGMP device:**

  A switch or router running IGMP traffic control features.

- **IGMP host:**

  An end-node device running an IGMP (multipoint or multicast communication) application.

- **Querier:**

  A required IGMP device that facilitates IGMP protocol and traffic flow on a given LAN. This device tracks which ports are connected to devices (IGMP clients) that belong to specific multicast groups and triggers updates of this information. A querier uses data received from the queries to determine whether to forward or block multicast traffic on specific ports. When the switch has an IP address on a given VLAN, the switch automatically operates as a querier for that VLAN if it does not detect a multicast router or another switch functioning as a querier. When enabled (the default state), the switch's querier function eliminates the need for a multicast router. In most cases, Hewlett Packard Enterprise recommends that you leave this parameter in the default enabled state even if you have a multicast router performing the querier function in your multicast group. For more information, see **How IGMP operates** on page 94.

IGMP operating features

In the factory default configuration, IGMP is disabled. If multiple VLANs are not configured, configure IGMP on the default VLAN (DEFAULT_VLAN; VID = 1); if multiple VLANs are configured, configure IGMP on a per-VLAN basis for every VLAN where this feature is needed.

With the CLI, you can also configure the following options:

- **Forward with high priority**: Disabling this parameter (the default) causes the switch or VLAN to process IP multicast traffic and other traffic, in the order received (usually normal priority). Enabling this parameter causes the switch or VLAN to give higher priority to IP multicast traffic than to other traffic.
- **Auto/blocked/forward**: You can configure individual ports to any of the following states:
    - **Auto (the default):** Causes the switch to interpret IGMP packets and to filter IP multicast traffic based on the IGMP packet information for ports belonging to a multicast group. Thus IGMP traffic is forwarded on a specific port only if an IGMP host or multicast router is connected to the port.
    - **Blocked:** Causes the switch to block IGMP joins arriving on the blocked port. A multicast stream will still flood out a blocked port if no active joins have been received.
    - **Forward:** Causes the switch to forward all IGMP and IP multicast transmissions through the port.
- **Operation with or without IP addressing:** Helps conserve IP addresses by enabling IGMP to run on VLANs that do not have an IP address. See **Operation with or without IP addressing** on page 95.
- **Querier capability:** The switch performs this function for IGMP on VLANs having an IP address when no other device in the VLAN is acting as querier. See **Using the switch as querier** on page 101.

NOTE: Whenever IGMP is enabled, the switch generates an Event Log message indicating whether querier functionality is enabled.

IP multicast traffic groups are identified by IP addresses in the range of 224.0.0.0 to 239.255.255.255 and incoming IGMP packets intended for reserved, or “well-known” multicast addresses automatically flood through all ports (except the port on which the packets entered the switch). For more on this topic, see **Excluding multicast addresses from IP multicast filtering** on page 101.

CLI: Configuring and displaying IGMP

Viewing the Current IGMP Configuration. The `show ip igmp config` command lists the IGMP configuration for all VLANs configured on the switch or for a specific VLAN.

Syntax:

```
show ip igmp config
```

IGMP configuration for all VLANs on the switch.

```
show ip igmp < vid > config
```

IGMP configuration for a specific VLAN on the switch, including per-port data.

(For IGMP operating status, see the appendix on monitoring and analyzing switch operation in the *ArubaOS-Switch Management and Configuration Guide*.)

For example, given the following VLAN and IGMP configurations on the switch:

VLAN ID	VLAN Name	IGMP Enabled	Forward with High Priority	Querier
1	DEFAULT_VLAN	Yes	No	No
22	VLAN-2	Yes	Yes	Yes
33	VLAN-3	No	No	No

The following examples display the data for `show ip igmp config`, `statistics`, and `group` commands:

IGMP configuration for all VLANs on a switch

```
Switch# show ip igmp config
 IGMP Service
  VLAN ID    VLAN NAME    IGMP Enabled Forward with High Priority Querier
  ---------- ------------ ------------ -------------------------- -------
  1          DEFAULT_VLAN Yes          No                         No
  22         VLAN-2       Yes          Yes                        Yes
  33         VLAN-3       No           No                         No
```

Displaying igmp high level statistics for all VLANs on a switch

Syntax: `show ip igmp statistics`

```
switch(config)# show ip igmp statistics
IGMP Service Statistic
  Total VLAN's with IGMP enabled:    33
  Current count of multicast groups joined:  21
IGMP Service Statistics
  VLAN ID VLAN Name        Total  Filtered Standard Static
  ------- ---------------- ------ -------- -------- -------
  1       DEFAULT_VLAN     52     50       0        2
  300     Office Client    80     75       5        0
  300     Data Center      1100   1000     99       1
```

Displaying igmp group address information

Syntax: `show ip igmp groups`

```
switch(config)# show ip igmp groups
IGMP Group Address Information
VLAN ID Group Address Expires         UpTime        Last Reporter  | Type
------- ------------- --------------- ------------- -------------- + -----
2       226.0.6.7     0h 2m 58s       1h 13m 4s     192.168.0.2    | Filter
2       226.0.6.8     0h 2m 58s       1h 13m 4s     192.168.0.2    | Standard
2       226.0.6.9     0h 2m 58s       1h 13m 4s     192.168.0.2    | Static
```

Displaying the IGMP configuration for a specific VLAN

The following **show ip igmp** command example shows the VLAN ID (**VID**) designation and the IGMP per-port configuration:

Figure 11: *Displaying the IGMP configuration for a specific VLAN*

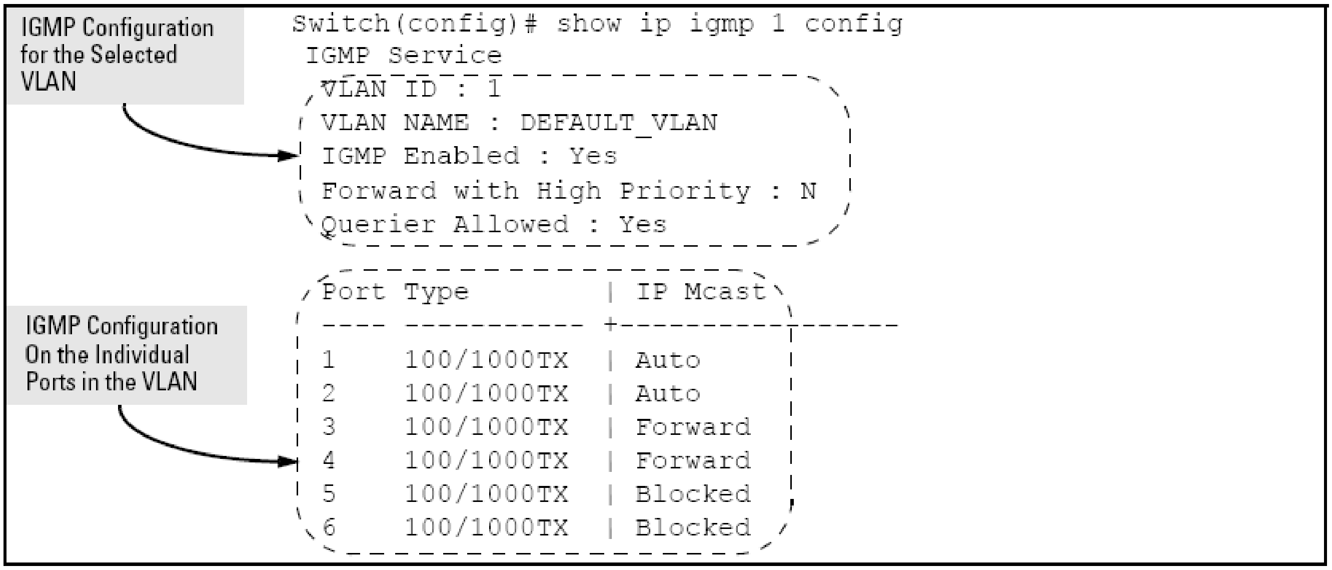

```
Switch(config)# show ip igmp 1 config
 IGMP Service
  VLAN ID : 1
  VLAN NAME : DEFAULT_VLAN
  IGMP Enabled : Yes
  Forward with High Priority : N
  Querier Allowed : Yes

  Port Type        | IP Mcast
  ---- ----------- +-----------------
  1    100/1000TX  | Auto
  2    100/1000TX  | Auto
  3    100/1000TX  | Forward
  4    100/1000TX  | Forward
  5    100/1000TX  | Blocked
  6    100/1000TX  | Blocked
```

IGMP Configuration for the Selected VLAN
IGMP Configuration On the Individual Ports in the VLAN

Enabling or disabling IGMP on a VLAN. You can enable IGMP on a VLAN with the last-saved or default IGMP configuration (whichever was most recently set) or you can disable IGMP on a selected VLAN.

NOTE: The `ip igmp` command must be executed in a VLAN context.

Syntax:

```
ip igmp
```

```
no ip igmp
```

Examples of enabling and disabling IGMP on the default VLAN (VID = 1):

Command syntax	Task
`# vlan 1 ip igmp`	Enables IGMP on VLAN 1.
`switch(vlan-1)# ip igmp`	Disables IGMP on VLAN 1.
`switch(config)# no vlan 1 ip igmp`	Disables IGMP on VLAN 1.

NOTE: If you disable IGMP on a VLAN and then later re-enable IGMP on that VLAN, the switch restores the last-saved IGMP configuration for that VLAN. For more information on switch memory operation, see the chapter on switch memory and configuration in the*ArubaOS-Switch Basic Operation Guide*.

You can also combine the `ip igmp` command with other IGMP-related commands, as described in the following sections.

Configuring Per-Port IGMP Packet Control.

Command syntax	Task
vlan < *vid* > ip igmp [auto *<port-list>* | blocked *<port-list>* | forward *<port-list>*]	Use this command in the VLAN context to specify how each port should handle IGMP traffic.
vlan < *vid* > ip igmp	Enables IGMP on the specified VLAN. In a VLAN context, use only ip igmp without the VLAN specifier.
vlan < *vid* > ip igmp auto *<port-list>* (default)	Filter multicast traffic on the specified ports. Forward IGMP traffic to hosts on the ports that belong to the multicast group for which the traffic is intended. (Also forward any multicast traffic through any of these ports that is connected to a multicast router.) This is the default IGMP port configuration.
vlan < *vid* > ip igmp blocked *<port-list>*	Drop all multicast traffic received from devices on the specified ports and prevent any outgoing multicast traffic from moving through these ports.
vlan < *vid* > ip igmp forward < port-list >	Forward all multicast traffic through the specified port.

For example, to configure IGMP as follows for VLAN 1 on ports 1 - 6:

- Ports 1 - 2: Auto
- Ports 3 - 4: Forward
- Ports 5 - 6: Block

Depending on privilege level, use the following commands to configure IGMP on VLAN 1:

```
switch(config)# vlan 1
switch(vlan-1)# ip igmp auto 1,2
switch(vlan-1)# ip igmp forward 3,4
switch(vlan-1)# ip igmp blocked 5,6
```

After executing the above commands, use the following command to display the VLAN and per-port configuration.

Configuring the querier function

The `ip igmp querier` command lets you disable or re-enable the ability for the switch to become querier on the specified VLAN. The default querier capability is “enabled”.

Syntax:

```
no vlan <vid> ip igmp querier
```

For example, the following `no vlan 1` command disables the querier function on VLAN 1.

```
switch(config)# no vlan 1 ip igmp querier
```

The following show command displays results of the previous querier command.

```
Switch# show ip igmp config
```

Web: Enabling and disabling IGMP

In the web browser, you can enable or disable IGMP per-VLAN. To configure other IGMP features, use the CLI on the switch console.

To enable or disable IGMP:

1. Click the **Configuration** tab.
2. Click the **Device Features** button.
3. If more than one VLAN is configured, use the VLAN pull-down menu to select the VLAN on which to enable or disable IGMP.
4. Use the Multicast Filtering (IGMP) menu to enable or disable IGMP.
5. Click the **Apply Changes** button to implement the configuration change.

For web-based help on how to use the web browser interface screen, click the **?** button on the web browser screen.

How IGMP operates

The Internet Group Management Protocol (IGMP) is an internal protocol of the Internet Protocol (IP) suite. IP manages multicast traffic by using switches, multicast routers and hosts that support IGMP. (In Hewlett Packard Enterprise's implementation of IGMP, a multicast router is not necessary as long as a switch is configured to support IGMP with the querier feature enabled.) A set of hosts, routers or switches that send or receive multicast data streams to or from the same sources is called a multicast group and all devices in the group use the same multicast group address.

Message types

The multicast group running IGMP uses three message types to communicate:

- **Query:**

  A message sent from the querier (multicast router or switch) asking for a response from each host belonging to the multicast group. If no multicast router supporting IGMP is present, then the switch assumes this function to elicit group membership information from the hosts on the network. (To disable the querier, use the CLI IGMP configuration MIB. See "Configuring the querier function" in **CLI: Configuring and displaying IGMP**)

- **Report (join):**

  A message sent by a host to the querier indicating that the host wants to be or is a member of a given group in the report message.

- **Leave group:**

  A message sent by a host to the querier indicating that the host has ceased to be a member of a specific multicast group.

IGMP multicasting

IGMP identifies members of a multicast group within a subnet and lets IGMP-configured hosts and routers join or leave multicast groups based on the following:

- An IP multicast packet includes the multicast group address to which the packet belongs.
- When an IGMP client connected to a switch port needs to receive multicast traffic from a specific group, it joins the group by sending an IGMP report (join request) to the network. The multicast group specified in the join request is determined by the requesting application running on the IGMP client.
- When a networking device with IGMP enabled receives the join request for a specific group, it forwards any IP multicast traffic it receives for that group through the port on which the join request was received.
- When the client is ready to leave the multicast group, it sends a Leave Group message to the network and ceases to be a group member.
- When the leave request is detected, the appropriate IGMP device ceases to transmit traffic for the designated multicast group through the port on which the leave request was received, as long as there are no other current members of that group on the affected port.

Displaying IGMP data

To display data showing active group addresses, reports, queries, querier access port and active group address data (port, type and access), see the appendix on monitoring and analyzing switch operation in the *ArubaOS-Switch Management and Configuration Guide*.

Supported standards and RFCs

The implementation of IGMP supports the following standards and operating capabilities:

- RFC2236 (IGMP V.2 with backwards support for IGMP V.1).
- IETF draft for IGMP and MLD snooping switches (for IGMP V1, V2 V3).
- Full IGMPv2 support and full support for IGMPv1 Joins.
- Ability to operate in IGMPv2 querier mode on VLANs with an IP address.

The implementation is subject to the following restrictions:

- Interoperability with RFC3376 (IGMPv3).
- Interoperability with IGMPv3 Joins. When the switch receives an IGMPv3 Join, it accepts the host request and begins forwarding the IGMP traffic. Thus ports that have not joined the group and are not connected to routers or the IGMP querier will not receive the group's multicast traffic.
- No support for the IGMPv3 “Exclude Source” or “Include Source” options in Join Reports; the group is simply joined from all sources.
- No support for becoming a version 3 querier. The switch becomes a version 2 querier in the absence of any other querier on the network.

NOTE:

IGMP is supported in the HPE MIB, not in standard IGMP MIBs, as the latter reduce Group Membership detail in switched environments.

Operation with or without IP addressing

You can configure IGMP on VLANs that do not have IP addressing. Using IGMP without IP addressing reduces the number of IP addresses you use and configure, significant in a network with many VLANs. The limitation on IGMP without IP addressing is that the switch cannot become querier on any VLANs for which it has no IP address; thus the network administrator must ensure that another IGMP device acts as querier. Hewlett Packard Enterprise also advises that an additional IGMP device be available as backup querier.

Table 10: *Comparison of IGMP operation with and without IP addressing*

IGMP Function available with IP Addressing configured on the VLAN	Available without IP Addressing?	Operating Differences without an IP Address
Forward multicast group traffic to any port on the VLAN that has received a join request for that multicast group.	Yes	None
Forward join requests (reports) to the querier.	Yes	None
Configure individual ports in the VLAN to **Auto** (the default) **Blocked**, or **Forward**.	Yes	None
Configure IGMP traffic forwarding to normal or high-priority forwarding.	Yes	None
Age-Out IGMP group addresses when the last IGMP client on a port in the VLAN leaves the group.	Yes	Requires that another IGMP device in the VLAN have an IP address and can operate as querier. This can be a multicast router or another switch configured for IGMP operation. Hewlett Packard Enterprise recommends that the VLAN also include a device operating as a backup querier in case the device operating as the primary querier fails.
Support Fast-Leave IGMP (below) and Forced Fast-Leave IGMP.	Yes	
Support automatic querier election.	No	Querier operation not available.
Operate as the querier.	No	Querier operation not available.
Available as a backup querier.	No	Querier operation not available.

Automatic Fast-Leave IGMP

IGMP Operation Presents a “Delayed Leave” Problem. Where multiple IGMP clients are connected to the same port on an IGMP device (switch or router), if only one IGMP client joins a given multicast group, then later sends a Leave Group message and ceases to belong to that group, the IGMP device retains that IGMP client in its IGMP table and continues forwarding IGMP traffic to the IGMP client until the querier triggers confirmation that no other group members exist on the same port. Thus the switch continues to transmit unnecessary multicast traffic through the port until the querier renews its multicast group status.

When unregistered multicasts are received on switches that support Data-Driven IGMP (“Smart” IGMP), the switch automatically drops them. Thus the sooner the IGMP Leave is processed, the sooner this multicast traffic stops flowing.

On switches that do not support Data-Driven IGMP, unregistered multicast groups are flooded to the VLAN rather than pruned. In this scenario, Fast-Leave IGMP can actually increase the problem of multicast flooding by removing the IGMP group filter before the querier has recognized the IGMP Leave. The querier continues to transmit the multicast group during this short time and because the group is no longer registered the switch then floods the multicast group to all ports. Because of such multicast flooding, the IGMP Fast-Leave feature is disabled by default on all switches that do not support Data-Driven IGMP. The feature can be enabled on these switches using an SNMP set of the following object:

```
hpSwitchIgmpPortForceLeaveState.< vid >.< port number>
```

However, this is not recommended as this will increase the amount of multicast flooding during the period between the client's IGMP Leave and the querier's processing of that Leave. For more on this topic, see the following figure.

Automatic Fast-Leave Operation The Fast-Leave operation applies if a switch port has the following characteristics:

- Is connected to only one end node.
- Is an IGMP client (the end node currently belongs to a multicast group).
- The end node subsequently leaves the multicast group.

The switch does not need to wait for the querier status update interval but instead immediately removes the IGMP client from its IGMP table and ceases transmitting IGMP traffic to the client. (If the switch detects multiple end nodes on the port, automatic Fast-Leave does not activate – regardless of whether one or more of these end nodes are IGMP clients.)

In the following figure, automatic Fast-Leave operates on the switch ports for IGMP clients “3A” and “5A”, but not on the switch port for IGMP clients “7A” and 7B, Server “7C” and printer “7D”.

Figure 12: *Automatic Fast-Leave IGMP Criteria*

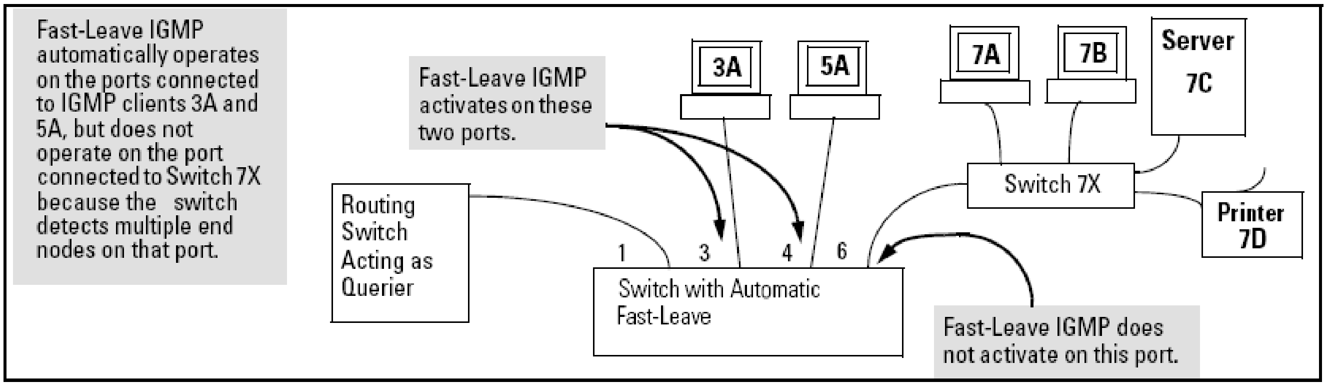
Fast-Leave IGMP automatically operates on the ports connected to IGMP clients 3A and 5A, but does not operate on the port connected to Switch 7X because the switch detects multiple end nodes on that port.
Fast-Leave IGMP activates on these two ports.
3A
5A
7A
7B
Server 7C
Switch 7X
Printer 7D
Routing Switch Acting as Querier
1
3
4
6
Switch with Automatic Fast-Leave
Fast-Leave IGMP does not activate on this port.

When client “3A” running IGMP is ready to leave the multicast group, it transmits a Leave Group message. Because the switch knows that there is only one end node on port 3, it removes the client from its IGMP table and halts multicast traffic (for that group) to port 3. If the switch is not the querier, it does not wait for the actual querier to verify that there are no other group members on port 3. If the switch itself is the querier, it does not query port 3 for the presence of other group members.

NOTE: Fast-Leave operation does not distinguish between end nodes on the same port that belong to different VLANs. Thus even if all devices on port 6 in in the preceding figure belong to different VLANs, Fast-Leave does not operate on port 6.

Using delayed group flush

This feature continues to filter IGMP-Left groups for a specified additional time. Delay in flushing the group filter prevents stale traffic from being forwarded by the server. Delayed group flush is enabled or disabled for the entire switch.

Syntax:

```
igmp delayed-flush <time period>
```

Enables the switch to continue to flush IGMP-Left groups for a specified period of time (0 - 255 seconds). The default setting is **Disabled**. To disable, reset the time period to zero.

Syntax:

```
show igmp delayed-flush
```

Displays the current setting for the switch.

Forced Fast-Leave IGMP

Forced Fast-Leave IGMP speeds up the process of blocking unnecessary IGMP traffic to a switch port that is connected to multiple end nodes. (This feature does not activate on ports where the switch detects only one end node.) For example, in **Figure 12: Automatic Fast-Leave IGMP Criteria** on page 97, even if you configured Forced Fast-Leave on all ports in the switch, the feature would activate only on port 6 (which has multiple end nodes) when a Leave Group request arrived on that port.

When a port having multiple end nodes receives a Leave Group request from one end node for a given multicast group “X”, Forced Fast-Leave activates and waits a short time to receive a join request from any other group “X” member on that port. If the port does not receive a join request for that group within the forced-leave interval, the switch then blocks any further group “X” traffic to the port.

Setting Fast-Leave and Forced Fast-Leave from the CLI

Previous, Fast-Leave and Forced Fast-Leave options for a port were set exclusively through the MIB. The following commands now allow a port to be configured for Fast-Leave or Forced Fast-leave operation from the CLI. These commands must be executed in a VLAN context.

Syntax:

```
ip igmp fastleave <port-list>
```

```
no ip igmp fastleave <port-list>
```

Enables IGMP Fast-Leaves on the specified ports in the VLAN (the default setting). In the Config context, use the VLAN specifier <*vid*>, for example, `vlan` < *vid* >`ip igmp fastleave` <*port-list*>. The `no` form disables Fast-Leave on the specified ports.

Syntax:

```
ip igmp forcedfastleave <port-list>
```

```
no ip igmp forcedfastleave <port-list>
```

Forces IGMP Fast-Leaves on the specified ports in the VLAN, even if they are cascaded.

To view the IGMP Fast-Leave status of a port use the `show running-config` or `show config` command.

Setting Forced Fast-Leave using the MIB

Fast-Leave and Forced Fast-Leave options for a port can also be set through the switch MIB (Management Information Base).

Table 11: *Forced Fast-Leave States*

Feature	Default	Settings	Function
Forced Fast-Leave state	**2** (disabled)	**1** (enabled)	Uses the **setmib** command to enable or disable Forced Fast-Leave on individual ports. When enabled on a port, Forced Fast-Leave operates only if the switch detects multiple end nodes (and at least one IGMP client) on that port
		2 (disabled)	

NOTE:

VLAN Numbers

In the switches covered in this manual, the **walkmib** and **setmib** commands use an internal VLAN number, not the VLAN ID or VID, to display or change many per-vlan features, such as the Forced Fast-Leave state. Because the internal VLAN number for the default VLAN is always 1, whether or not VLANs are enabled on the switch, examples herein use the default VLAN.

Listing the MIB-Enabled Forced Fast-Leave configuration

Forced Fast-Leave configuration data available in the switch MIB includes the state (enabled or disabled) for each port and the Forced-Leave Interval for all ports on the switch.

To List the Forced Fast-Leave State for all Ports in the Switch. In the CLI, use the **walkmib** command, as shown below.

Enter either of the following `walkmib` commands (generic or explicit):

`walkmib hpSwitchIgmpPortForcedLeaveState` (generic command)

OR

`walkmib 1.3.6.1.4.1.11.2.14.11.5.1.7.1.15.3.1.5` (explicit command)

The result shows the Forced Fast-Leave state for all ports in the switch, by VLAN. (A port belonging to more than one VLAN is shown once for each VLAN; if multiple VLANs are not configured, all ports are shown as members of the default VLAN.) For example, the following figure shows output of the `walkmib` command.

Figure 13: *Forced Fast-Leave output where all ports are members of the default VLAN*

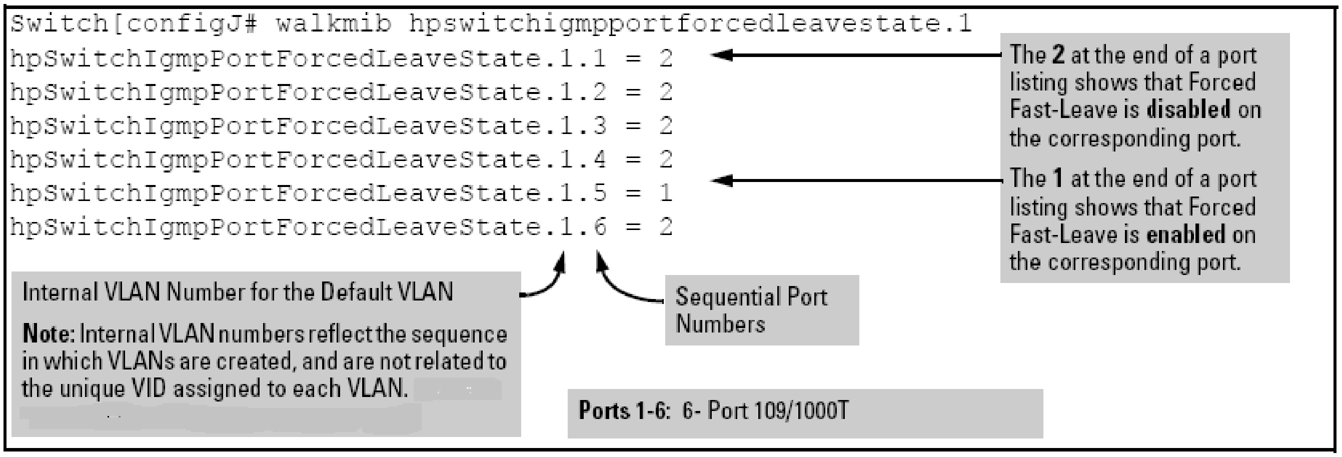

```
Switch[configJ# walkmib hpswitchigmpportforcedleavestate.1
hpSwitchIgmpPortForcedLeaveState.1.1 = 2
hpSwitchIgmpPortForcedLeaveState.1.2 = 2
hpSwitchIgmpPortForcedLeaveState.1.3 = 2
hpSwitchIgmpPortForcedLeaveState.1.4 = 2
hpSwitchIgmpPortForcedLeaveState.1.5 = 1
hpSwitchIgmpPortForcedLeaveState.1.6 = 2
```

The **2** at the end of a port listing shows that Forced Fast-Leave is **disabled** on the corresponding port.

The **1** at the end of a port listing shows that Forced Fast-Leave is **enabled** on the corresponding port.

Internal VLAN Number for the Default VLAN

Note: Internal VLAN numbers reflect the sequence in which VLANs are created, and are not related to the unique VID assigned to each VLAN.

Sequential Port Numbers

Ports 1-6: 6- Port 109/1000T

To show the Forced Fast-Leave state for a single port

Use the following `getmib` command (see the following figure).

Syntax:

```
getmib hpSwitchIgmpPortForcedLeaveState.<vlan number><.port number>
```

OR

```
getmib 1.3.6.1.4.1.11.2.14.11.5.1.7.1.15.3.1.5.<vlan number><.port number>
```

For example, the following `getmib` command shows the state for port 6 on the default VLAN.

Figure 14: *Forced Fast-Leave state for a single port on the default VLAN*

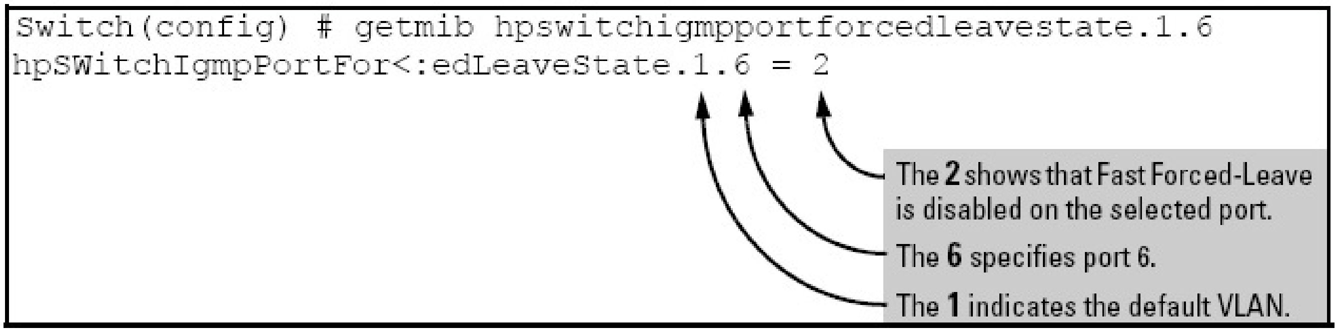

```
Switch(config) # getmib hpswitchigmpportforcedleavestate.1.6
hpSWitchIgmpPortFor<:edLeaveState.1.6 = 2
```

The **2** shows that Fast Forced-Leave is disabled on the selected port.

The **6** specifies port 6.

The **1** indicates the default VLAN.

Configuring per-port Forced Fast-Leave IGMP

In the factory-default configuration, Forced Fast-Leave is disabled for all ports on the switch. To enable (or disable) this feature on individual ports, use the switch **setmib** command.

Configuring Per-Port Forced Fast-Leave IGMP on Ports. This procedure enables or disables Forced Fast-Leave on ports in a given VLAN.

```
switch(config)# setmib hpswitchigmpportforcedleavestate.1.6 -i 1
hpSwitchIgmpPortForcedLeaveState.1.6 = 2
```

where **1** in .1.6 is the default VLAN, **6** in .1.6 indicates port 6 and **= 2** verifies Forced Fast-Leave disabled.

Syntax:

```
setmib hpSwitchIgmpPortForcedLeaveState.< vlan number >< .port number >-i < 1 | 2 >
```

OR

```
setmib 1.3.6.1.4.1.11.2.14.11.5.1.7.1.15.3.1.5.< vlan number >< .port number > -i <
1 | 2 >
```

where:

Table 12: *Forced Fast-Leave values*

1	Enabled
2	Disabled

For example, suppose that your switch has six ports as members of the default VLAN. To enable Forced Fast-Leave on port 6, you would execute the following command to obtain the result.

Changing the Forced Fast-Leave Configuration on Port 6.

```
switch(config)# setmib hpswitchigmpportforcedleavestate.1.6 -i 1
hpSwitchIgmpPortForcedLeaveState.1.6 = 1
```

where **1** in .1.6 is the default VLAN, **6** in .1.6 indicates port 6 and **= 1** verifies Forced Fast-Leave enabled.

Using the switch as querier

Querier operation

The function of the IGMP querier is to poll other IGMP-enabled devices in an IGMP-enabled VLAN to elicit group membership information. The switch performs this function if there is no other device in the VLAN, such as a multicast router, to act as querier. Although the switch automatically ceases querier operation in an IGMP-enabled VLAN if it detects another querier on the VLAN, you can also use the CLI to disable the querier capability for that VLAN.

NOTE:

A querier is required for proper IGMP operation. Thus, if you disable the querier function on a switch, ensure that there is an IGMP querier (and, preferably, a backup querier) available on the same VLAN.

If the switch becomes the querier for a particular VLAN (for example, the DEFAULT_VLAN), then subsequently detects queries transmitted from another device on the same VLAN, the switch ceases to operate as the querier for that VLAN. If this occurs, the switch Event Log lists a pair of messages such as:

```
I 01/15/01 09:01:13 igmp: DEFAULT_VLAN: Other Querier detected
```

```
I 01/15/01 09:01:13 igmp: DEFAULT_VLAN: This switch is no longer Querier
```

In the above scenario, if the other device ceases to operate as a querier on the default VLAN, then the switch detects this change and can become the querier as long as it is not preempted by some other IGMP querier on the VLAN. In this case, the switch Event Log lists messages such as the following to indicate that the switch has become the querier on the VLAN:

```
I 01/15/01 09:21:55 igmp: DEFAULT_VLAN: Querier Election in process
```

```
I 01/15/01 09:22:00 igmp: DEFAULT_VLAN: This switch has been elected as Querier
```

Excluding multicast addresses from IP multicast filtering

Each multicast host group is identified by a single IP address in the range 224.0.0.0 through 239.255.255.255. Specific groups of consecutive addresses in this range are “well-known” addresses and are reserved for predefined host groups. IGMP does **not** filter these addresses, so any packets the switch receives for such addresses are flooded out all ports assigned to the VLAN on which they were received (except the port on which the packets entered the VLAN).

The following table lists the 32 well-known address groups (8192 total addresses) that IGMP does not filter on the switches covered in this manual.

Table 13: *IP multicast address groups excluded from IGMP filtering*

In the table, x is any value from 0 to 255.

Groups of consecutive addresses in the range of 224.0.0.x to 239.0.0.x		Groups of consecutive addresses in the range of 224.128.0.x to 239.128.0.x	
224.0.0.x	232.0.0.x	224.128.0.x	232.128.0.x
225.0.0.x	233.0.0.x	225.128.0.x	233.128.0.x
226.0.0.x	234.0.0.x	226.128.0.x	234.128.0.x
227.0.0.x	235.0.0.x	227.128.0.x	235.128.0.x
228.0.0.x	236.0.0.x	228.128.0.x	236.128.0.x
229.0.0.x	237.0.0.x	229.128.0.x	237.128.0.x
230.0.0.x	238.0.0.x	230.128.0.x	238.128.0.x
231.0.0.x	239.0.0.x	231.128.0.x	239.128.0.x

NOTE: IP multicast addresses occur in the range 224.0.0.0 through 239.255.255.255 (corresponding to the Ethernet multicast address range 01005e-000000 through 01005e-7fffff). Where a switch has a static Traffic/Security filter configured with a "Multicast" filter type and a "Multicast Address" in this range, the switch uses the static filter unless IGMP learns of a multicast group destination in this range. In this case, IGMP dynamically takes over the filtering function for the multicast destination addresses) for as long as the IGMP group is active. If the IGMP group subsequently deactivates, the switch returns filtering control to the static filter.

Reserved addresses excluded from IP multicast (IGMP) filtering.

Traffic to IP multicast groups in the IP address range 224.0.0.0 to 224.0.0.255 are always flooded because addresses in this range are well known or reserved addresses. Thus, if IP multicast is enabled and there is an IP multicast group within the reserved address range, traffic to that group is flooded, not filtered, by the switch.

Number of IP multicast addresses allowed. Multicast filters and IGMP filters (addresses) together can total up to 255 in the switch. If multiple VLANs are configured, then each filter is counted once per VLAN in which it is used.

Overview of MSTP

Without spanning tree, having more than one active path between a pair of nodes causes loops in the network, which can result in duplication of messages leading to a "broadcast storm" that can bring down the network.

NOTE: MSTP cannot protect against loops when there is an unmanaged device on the network that drops spanning tree packets, or may fail to detect loops where this is an edge port configured with client authentication (802.1X, Web and MAC authentication). To protect against the formation of loops in these cases, you can use the loop protection feature.

Multiple-Instance spanning tree operation (802.1s) ensures that only one active path exists between any two nodes in a spanning tree instance. A spanning tree instance comprises a unique set of VLANs, and belongs to a specific spanning tree region. A region can comprise multiple spanning tree instances (each with a different set of VLANs), and allows one active path among regions in a network. Applying VLAN tagging to the ports in a multiple-instance spanning tree network enables blocking of redundant links in one instance while allowing forwarding over the same links for non-redundant use by another instance.

VLAN/Instance groupings

Suppose that there are three switches in a region configured with VLANs grouped into two instances, as follows:

VLANs	Instance 1	Instance 2
10, 11, 12	Yes	No
20, 21, 22	No	Yes

The logical and physical topologies resulting from these VLAN/Instance groupings result in blocking on different links for different VLANs:

A multiple spanning tree application

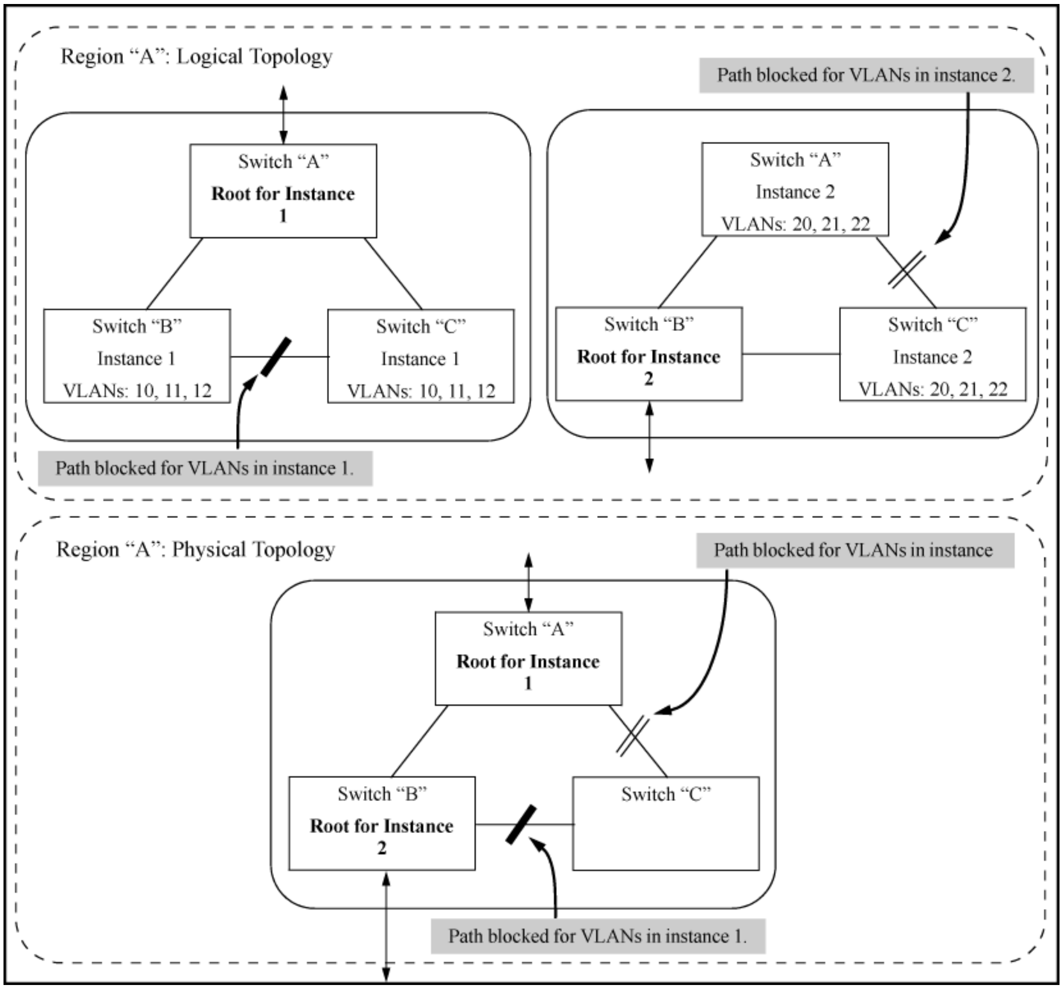

MSTP structure

MSTP maps active, separate paths through separate spanning tree instances and between MST regions. Each MST region comprises one or more MSTP switches. Note that MSTP recognizes an STP or RSTP LAN as a distinct spanning tree region.

Figure 15: *An MSTP network with legacy STP and RSTP devices connected*

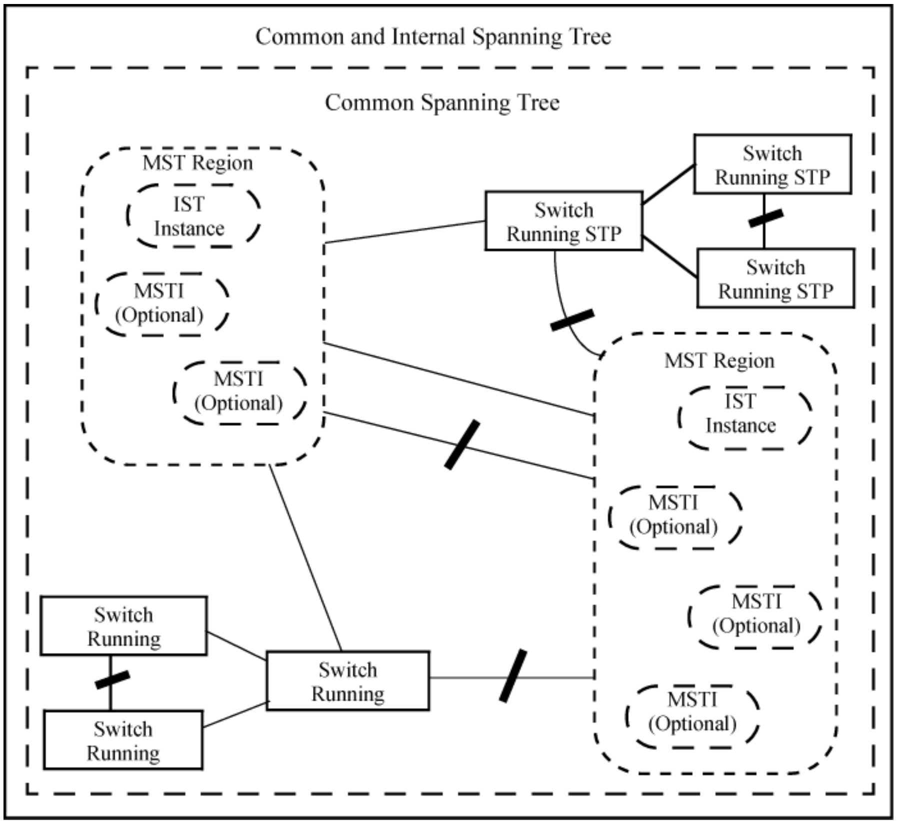

How MSTP operates

In the factory default configuration, spanning tree operation is off. Also, the switch retains its currently configured spanning tree parameter settings when disabled. Thus, if you disable spanning tree, then later re-enable it, the parameter settings will be the same as before spanning tree was disabled. The switch also includes a `pending` feature that enables you to exchange MSTP configurations with a single command.

NOTE:

The switch automatically senses port identity and type, and automatically defines spanning tree parameters for each type, and parameters that apply across the switch. Although these parameters can be adjusted, HPE strongly recommends leaving these settings in their default configurations unless the proposed changes have been supplied by an experienced network administrator who has a strong understanding of the IEEE 802.1D/w/s standards and operation.

802.1s Multiple Spanning Tree Protocol (MSTP)

The switches covered in this guide use the IEEE 802.1s Multiple Spanning Tree Protocol (MSTP) standard.

The 802.1D and 802.1w spanning tree protocols operate without regard to a network's VLAN configuration, and maintain one common spanning tree throughout a bridged network. Thus, these protocols map one loop-free, logical topology on a given physical topology. The 802.1s Multiple Spanning Tree protocol (MSTP) uses VLANs to create multiple spanning trees in a network, which significantly improves network resource utilization while maintaining a loop-free environment.

While the per-VLAN spanning tree approach adopted by some vendors overcomes the network utilization problems inherent in using STP or RSTP, using a per-VLAN technology with multiple VLANs can overload the switch's CPU. MSTP on the switches covered in this guide complies with the IEEE 802.1s standard, and extends STP and RSTP functionality to map multiple independent spanning tree instances onto a physical topology. With MSTP, each spanning tree instance can include one or more VLANs and applies a separate, per-instance forwarding topology. Thus, where a port belongs to multiple VLANs, it may be dynamically blocked in one spanning tree instance, but forwarding in another instance. This achieves load-balancing across the network while keeping the switch's CPU load at a moderate level (by aggregating multiple VLANs in a single spanning tree instance). MSTP provides fault tolerance through rapid, automatic reconfiguration if there is a failure in a network's physical topology.

With MSTP-capable switches, you can create a number of MST regions containing multiple spanning tree instances. This requires the configuration of a number of MSTP-capable switches. However, it is not necessary to do this. You can enable MSTP on an MSTP-capable switch and a spanning tree instance is created automatically. This instance always exists by default when spanning tree is enabled, and is the spanning tree instance that communicates with STP and RSTP environments. The MSTP configuration commands operate exactly like RSTP commands and MSTP is backward-compatible with the RSTP-enabled and STP-enabled switches in your network.

CAUTION:

Spanning tree interprets a switch mesh as a single link. Because the switch automatically gives faster links a higher priority, the default MSTP parameter settings are usually adequate for spanning tree operation. Because incorrect MSTP settings can adversely affect network performance, do not change the MSTP settings from their default values unless you have a strong understanding of how spanning tree operates.

In a mesh environment, the default MSTP timer settings (`Hello Time` and `Forward Delay`) are usually adequate for MSTP operation. Because a packet crossing a mesh may traverse several links within the mesh, using smaller-than-default settings for the MSTP `Hello Time` and `Forward Delay` timers can cause unnecessary topology changes and end-node connectivity problems.

MST regions

All MSTP switches in a given region must be configured with the same VLANs, and each MSTP switch within the same region must have the same VLAN-to-instance assignments. In addition, a VLAN can belong to only one instance within any region. Within a region:

- All of the VLANs belonging to a given instance compose a single, active spanning tree topology for that instance.
- Each instance operates independently of other regions.

Between regions, there is a single, active spanning tree topology.

How separate instances affect MSTP

Assigning different groups of VLANs to different instances ensures that those VLAN groups use independent forwarding paths. For example, in the following figure, each instance has a different forwarding path.

Figure 16: *Active topologies built by three independent MST instances*

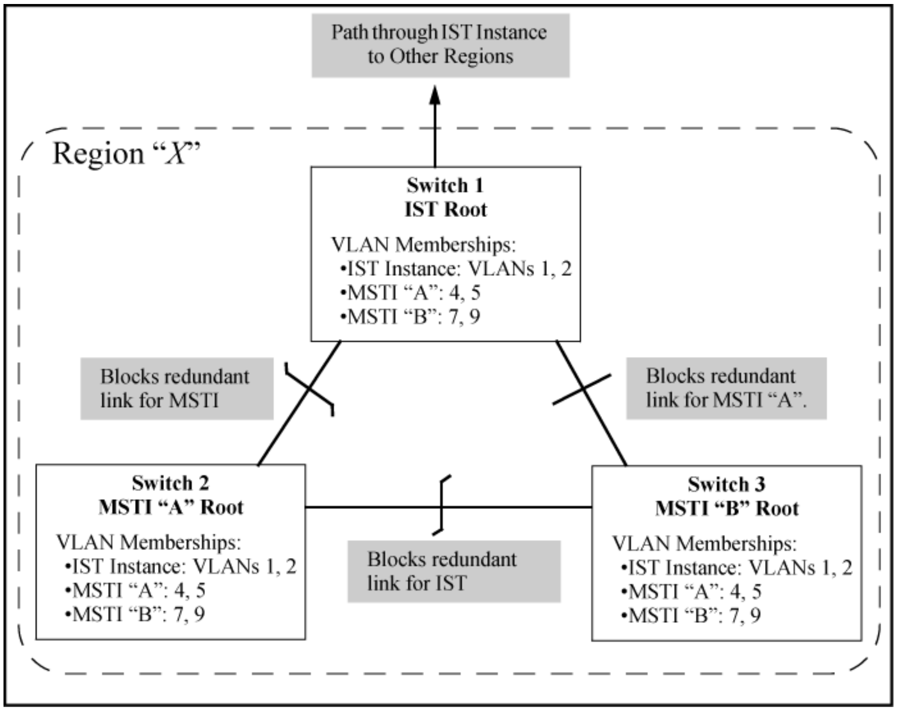
Path through IST Instance to Other Regions
Region "X"
Switch 1
IST Root
VLAN Memberships:
•IST Instance: VLANs 1, 2
•MSTI "A": 4, 5
•MSTI "B": 7, 9
Blocks redundant link for MSTI
Blocks redundant link for MSTI "A".
Switch 2
MSTI "A" Root
VLAN Memberships:
•IST Instance: VLANs 1, 2
•MSTI "A": 4, 5
•MSTI "B": 7, 9
Blocks redundant link for IST
Switch 3
MSTI "B" Root
VLAN Memberships:
•IST Instance: VLANs 1, 2
•MSTI "A": 4, 5
•MSTI "B": 7, 9

While allowing only one active path through a given instance, MSTP retains any redundant physical paths in the instance to serve as backups (blocked) paths in case the existing active path fails. Thus, if an active path in an instance fails, MSTP automatically activates (unblocks) an available backup to serve as the new active path through the instance for as long as the original active path is down. Note also that a given port may simultaneously operate in different states (forwarding or blocking) for different spanning tree instances within the same region. This depends on the VLAN memberships to which the port is assigned. For example, if a port belongs to VLAN 1 in the IST instance of a region and also belongs to VLAN 4 in MSTI "x" in the same region, the port may apply different states to traffic for these two different instances.

Within a region, traffic routed between VLANs in separate instances can take only one physical path. To ensure that traffic in all VLANs within a region can travel between regions, all of the boundary ports for each region should belong to all VLANs configured in the region. Otherwise, traffic from some areas within a region could be blocked from moving to other regions.

All MSTP switches (including STP and RSTP switches) in a network use BPDUs (Bridge Protocol Data Units) to exchange information from which to build multiple, active topologies in the individual instances within a region and between regions. From this information:

- The MSTP switches in each LAN segment determine a designated bridge and designated port or trunk for the segment.
- The MSTP switches belonging to a particular instance determine the root bridge and root port or trunk for the instance.
- For the IST instance within a region, the MSTP switches linking that region to other regions (or to STP or RSTP switches) determine the IST root bridge and IST root port or trunk for the region. (For any Multiple

spanning tree instance—MSTI—in a region, the regional root may be a different switch that is not necessarily connected to another region.)

- The MSTP switches block redundant links within each LAN segment, across all instances, and between regions, to prevent any traffic loops.

As a result, each individual instance (spanning tree) within a region determines its regional root bridge, designated bridges, and designated ports or trunks.

Regions, legacy STP and RSTP switches, and the Common Spanning Tree (CST)

The IST instance and any MST instances in a region exist only within that region. Where a link crosses a boundary between regions (or between a region and a legacy STP or RSTP switch), traffic is forwarded or blocked as determined by the Common Spanning Tree (CST). The CST ensures that there is only one active path between any two regions, or between a region and a switch running STP and RSTP. (See the figure in **MSTP structure** on page 105.)

MSTP operation with 802.1Q VLANs

As indicated in the preceding sections, within a given MST instance, a single spanning tree is configured for all VLANs included in that instance. This means that if redundant physical links exist in separate VLANs within the same instance, MSTP blocks all but one of those links. However, you can prevent the bandwidth loss caused by blocked redundant links for different VLANs in an instance by using a port trunk. The following example shows how you can use a port trunk with 802.1Q (tagged) VLANs and MSTP without unnecessarily blocking any links or losing any bandwidth.

Using a trunked link to support multiple VLAN connectivity within the same MST instance

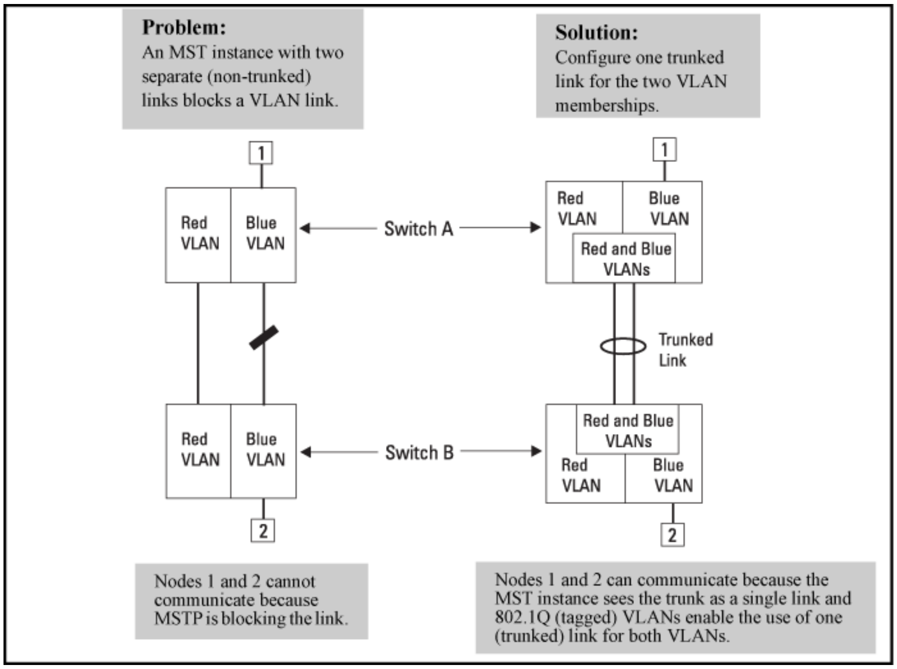
Problem:
An MST instance with two separate (non-trunked) links blocks a VLAN link.
Solution:
Configure one trunked link for the two VLAN memberships.
1
Red VLAN
Blue VLAN
Switch A
Red VLAN
Blue VLAN
Red and Blue VLANs
Trunked Link
Red VLAN
Blue VLAN
Switch B
Red and Blue VLANs
Red VLAN
Blue VLAN
2
Nodes 1 and 2 cannot communicate because MSTP is blocking the link.
Nodes 1 and 2 can communicate because the MST instance sees the trunk as a single link and 802.1Q (tagged) VLANs enable the use of one (trunked) link for both VLANs.

NOTE:

All switches in a region should be configured with the VLANs used in that region, and all ports linking MSTP switches together should be members of all VLANs in the region. Otherwise, the path to the root for a given VLAN will be broken if MSTP selects a spanning tree through a link that does not include that VLAN.

MSTP compatibility with RSTP or STP

IEEE 802.1s MSTP includes RSTP functionality and is designed to be compatible with both IEEE 802.1D and 802.1w spanning tree protocols. Using the default configuration values, your switches will interoperate effectively with RSTP and STP devices. MSTP automatically detects when the switch ports are connected to non-MSTP devices in the spanning tree and communicates with those devices using 802.1D or 802.1w STP BPDU packets, as appropriate.

To enable effective interoperation with STP (802.1D) configured devices, however, you may need to adjust the default configuration values. Here are two such examples:

- The rapid state transitions employed by MSTP may result in an increase in the rates of frame duplication and misordering in the switched LAN. To allow the switch to support applications and protocols that may be sensitive to frame duplication and misordering, you can disable rapid transitions by setting the Force Protocol Version parameter to STP-compatible. The value of this parameter applies to all ports on the switch.
- One of the benefits of MSTP is the implementation of a larger range of port path costs, which accommodates higher network speeds. However, this can create some incompatibility between devices running the older 802.1D STP. You can adjust to this incompatibility by implementing the global spanning tree legacy-path cost command.

RSTP and MSTP implement a greater range of path costs than 802.1D STP, and use different default path cost values to account for higher network speeds. These values are shown in the following table.

Port type	802.1D STP path cost	RSTP and MSTP path cost
10 Mbps	100	2 000 000
100 Mbps	10	200 000
1 Gbps	5	20 000

Because the maximum value for the path cost allowed by 802.1D STP is 65535, devices running that version of spanning tree cannot be configured to match the values defined by MSTP, at least for 10 Mbps and 100 Mbps ports. In LANs where there is a mix of devices running 802.1D STP, RSTP, and MSTPs, you should reconfigure the devices so the path costs match for ports with the same network speeds.

Preconfiguring an MSTP regional topology

The MSTP VLAN configuration enhancement allows you to preconfigure an MSTP regional topology and ensure that the same VLAN ID-to-MSTI assignments exist on each MSTP switch in the region.

CAUTION: When this software version is installed, the prior VLAN ID-to-MSTI mappings do not change. However, this enhancement is not backward-compatible. If you install a software version earlier than this version, and you have configured MSTI entries instances mapped to VLANs, they will be removed from the configuration file when booting to the prior version of software. Do one of the following to install or reload a prior version of the software:

- Remove all MSTP mappings from the configuration file, then reconfigure the instance mapping after running the desired software version.
- Save the current configuration file before updating the software to a new version. If you later reload this older version of the software, use this configuration file when you reload the older version.

The default behavior of the `spanning-tree instance vlan` command changes so that, before a static VLAN is configured or a dynamic VLAN is learned on the switch, you can preconfigure its VLAN ID-to-MSTI mapping. Later, when the VLAN is created, it is automatically assigned to the MSTI to which it was previously mapped.

By supporting preconfigured VLAN ID-to-MSTI topologies, the VLAN configuration enhancement provides the following benefits:

- Scalability: In a network design in which you plan to use a large number of VLANs, you can preconfigure identical VLAN ID-to-MSTI mappings on all switches in a single, campus-wide MST region, regardless of the specific VLANs that you later configure on each switch. After the initial VLAN ID-to-MSTI mapping, you can decide on the exact VLANs that you need on each switch.All switches in a region must be configured with the same VLAN ID-to-MSTI mappings and the same MSTP configuration identifiers (region name and revision number).
- Flexibility: By preconfiguring identical VLAN ID-to-MSTI mappings on all switches in an MST region, you can combine switches that support different maximum numbers of VLANs.
- Network stability: You can reduce the interruptions in network connectivity caused by the regeneration of spanning trees in the entire network each time a configuration change in VLAN-to-MSTI mapping is detected on a switch. The negative impact on network performance is reduced if all newly created VLANs are pre-mapped to the correct MST instances. Later, VLAN creation and deletion are ignored by MSTP and no interruption in spanning tree traffic occurs.
- Usability: Dynamically learned GVRP VLANs can be mapped to MSTIs and support MSTP load balancing.

Preconfiguring VLANs in an MST instance

When configuring an MSTP regional topology, multiple spanning tree instances are created. Each MST instance provides a fully connected active topology for a particular set of VLANs.

Each switch in an MSTP region is configured with the following set of common parameters:

- Region name (`spanning-tree config-name`)
- Region revision number (`spanning-tree config-revision`)
- Identical VLAN ID-to-MSTI mapping (`spanning-tree instance vlan`)

Syntax:

```
spanning-tree instance 1..16 vlan vid [vid..vid]
```

```
no spanning-tree instance 1..16 vlan vid [vid..vid]
```

Configuring MSTP on the switch automatically configures the IST instance and places all statically and dynamically configured VLANs on the switch into the IST instance. This command creates a new MST instance (MSTI) and moves the VLANs specified from the IST to the MSTI.

You must map at least one VLAN to an MSTI when you create it. You cannot map a VLAN ID to more than one instance. You can create up to 16 MSTIs in a region.

The `no` form of the command removes one or more VLANs from the specified MSTI. If no VLANs are specified, the `no` form of the command deletes the specified MSTI.

When removing a VLAN from an MSTI, the VLAN returns to the IST instance, where it remains or is re-assigned to another MSTI configured in the region.

NOTE: The valid VLAN IDs to map to a specified MSTI are from 1 to 4094. The VLAN ID-to-MSTI mapping does not require a VLAN to be already configured on the switch. The MSTP VLAN enhancement allows preconfiguring MSTP topologies before the VLAN IDs associated with each instance exist on a switch.

When using preconfigured VLAN ID-to-MSTI topologies, ensure that MSTP switches remain in the same region by mapping all VLAN IDs used in the region to the same MSTIs on each regional switch.

Each MST instance supports a different set of VLANs. A VLAN that is mapped to an MST instance cannot be a member of another MST instance.

The MSTP VLAN configuration enhancement allows you to ensure that the same VLAN ID-to-MSTI assignments exist on each MSTP switch in a region. Before a static VLAN is configured or a dynamic VLAN is learned on the switch, use the `spanning-tree instance vlan` command to map VLANs to each MST instance in the region. Later, when the VLAN is created, the switch automatically assigns it to the MST instance to which you had previously mapped it.

Configuring MSTP instances with the VLAN range option (Example)

Using the `spanning-tree instance` command with the VLAN range option configures the entire range of VLANs, even if the range includes VLANs that are not currently present on the switch.

Mapping VLANs to MSTP Instance

If VLANs 1, 5, and 7 are currently present and you enter the following command, all the VLANs from 1 through 10 are included, even those VLANs that are not present.

```
switch(config)# spanning-tree instance 1 vlan 1-10
```

On switches other than those covered by this guide, only the VLANs that are present will be included, that is, only VLANs 1, 5, and 7. The switch will map these VLANs to MSTP Instance 1, which results in a Configuration Digest that is not the same as the Configuration Digest for the switches running this enhancement.

Below, the example shows an MSTP instance configured with the VLAN range option. All the VLANs are included in the instance whether they exist or not.

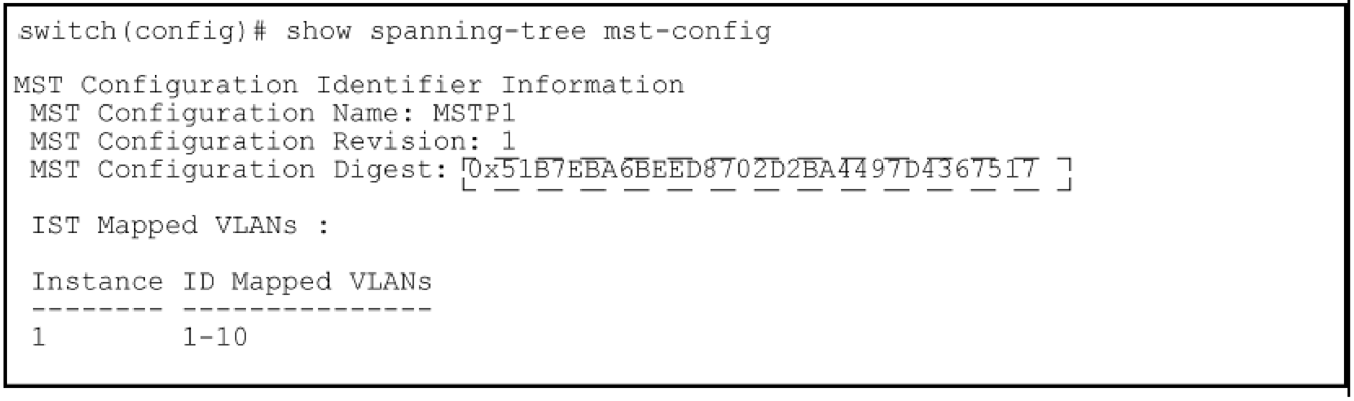

```
 switch(config)# show spanning-tree mst-config

MST Configuration Identifier Information
 MST Configuration Name: MSTP1
 MST Configuration Revision: 1
 MST Configuration Digest: 0x51B7EBA6BEED8702D2BA4497D4367517

 IST Mapped VLANs :

 Instance ID Mapped VLANs
 -------- ---------------
 1        1-10
```

Configuration Digest value

The Configuration Digest value shown below is not the same as in the above example indicating that these switches do not operate in the same instance.

The Common Spanning Tree (CST) will still have the correct root associations.

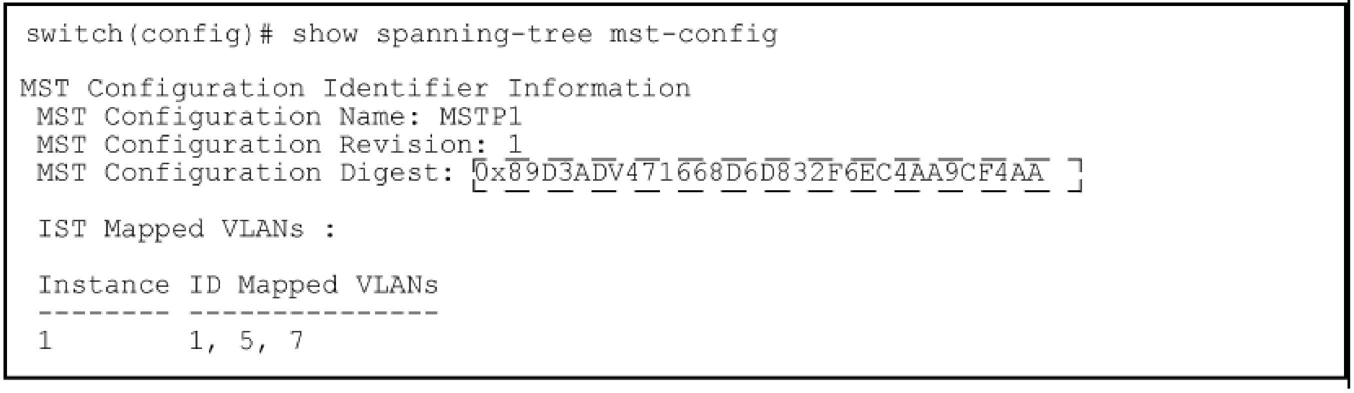

```
switch(config)# show spanning-tree mst-config

MST Configuration Identifier Information
 MST Configuration Name: MSTP1
 MST Configuration Revision: 1
 MST Configuration Digest: 0x89D3ADV471668D6D832F6EC4AA9CF4AA

 IST Mapped VLANs :

 Instance ID Mapped VLANs
 -------- ---------------
 1        1, 5, 7
```

Saving the current configuration before a software upgrade

Before updating to a new version of software, follow these steps:

Procedure

1. Enter the `show config files` command to display your current configuration files:

   ```
   switch(config)# show config files

   Configuration files:

    id | act pri sec | name
    ---+-------------+--------------------
     1 |  *   *   *  | config1
     2 |             | config2
     3 |             |
   ```

2. To save a configuration file for software version K.12.43, for example, type:

   ```
   switch(config)# copy config config1 config configK1243.cfg
   ```

   Choose any name for the saved configuration file that you prefer.

3. Display the configuration files as shown in the following example. Note the newly created configuration file listed.

   ```
   switch(config)# show config files

   Configuration files:

    id | act pri sec | name
    ---+-------------+----------------------
     1 |  *   *   *  | config1
     2 |             | config2
     3 |             | configK1243.cfg
   ```

4. Update the switch to the desired version, for example, K.12.51. Enter the `show flash` command to see the results. The switch is now running the software version K.12.51.

```
switch(config)# show flash

Image           Size(Bytes)   Date   Version  Build #:
-----           ----------- -------- -------  -------
Primary Image   : 6771179   04/17/08 K.12.51    304
Secondary Image : 7408949   11/06/08 K.12.43    123
Boot Rom Version: K.12.12
Default Boot    : Primary
```

5. To run the prior software version (K.12.43 in this example), type:

```
switch(config)# boot system flash secondary config configK1243.cfg
```

6. After rebooting, the switch is running software version K.12.43 and is using the configuration file that you saved for this software version, configK1243.cfg.
7. You can also save the K.12.43 configuration file on a TFTP server. To reload the K.12.43 version of the software again, reload the configuration file before doing the reload.

Types of Multiple Spanning Tree Instances

A multiple spanning tree network comprises separate spanning tree instances existing in an MST region. (There can be multiple regions in a network.) Each instance defines a single forwarding topology for an exclusive set of VLANs. By contrast, an STP or RSTP network has only one spanning tree instance for the entire network, and includes all VLANs in the network. (An STP or RSTP network operates as a single-instance network.) A region can include two types of STP instances:

- Internal spanning tree Instance (IST Instance) This is the default spanning tree instance in any MST region. It provides the root switch for the region and comprises all VLANs configured on the switches in the region that are not specifically assigned to Multiple Spanning Tree Instances (MSTIs, described below).Within a region, the IST instance provides a loop-free forwarding path for all VLANs associated with it. VLANs that are not associated with an MSTI are, by default, associated with the IST instance. Note that the switch automatically places dynamic VLANs (resulting from GVRP operation) in the IST instance. Dynamic VLANs cannot exist in an MSTI (described below).
- Multiple Spanning Tree Instance (MSTI)This type of configurable spanning tree instance comprises all static VLANs you specifically assign to it, and must include at least one VLAN. The VLANs you assign to an MSTI must initially exist in the IST instance of the same MST region. When you assign a static VLAN to an MSTI, the switch removes the VLAN from the IST instance. (Thus, you can assign a VLAN to only one MSTI in a given region.) All VLANs in an MSTI operate as part of the same single spanning tree topology. (The switch does not allow dynamic VLANs in an MSTI.)

CAUTION:

When you enable MSTP on the switch, the default MSTP spanning tree configuration settings comply with the values recommended in the IEEE 802.1s Multiple Spanning Tree Protocol (MSTP) standard. Inappropriate changes to these settings can result in severely degraded network performance. For this reason, HPE strongly recommends that changing these default settings be reserved only for experienced network administrators who have a strong understanding of the IEEE 802.1D/w/s standards and operation.

Planning an MSTP application

Before configuring MSTP, keep in mind the following tips and considerations:

- Ensure that the VLAN configuration in your network supports all the forwarding paths necessary for the desired connectivity. All ports connecting one switch to another within a region and one switch to another between regions should be configured as members of all VLANs configured in the region.
- Configure all ports or trunks connecting one switch to another within a region as members of all VLANs in the region. Otherwise, some VLANs could be blocked from access to the spanning tree root for an instance or for the region.
- Plan individual regions based on VLAN groupings. That is, plan on all MSTP switches in a given region supporting the same set of VLANs. Within each region, determine the VLAN membership for each spanning tree instance. (Each instance represents a single forwarding path for all VLANs in that instance.)
- Verify that there is one logical spanning tree path through the following:
  - Any interregional links
  - Any IST (Internal Spanning Tree) or Multiple Spanning Tree Instance within a region
  - Any legacy (802.1D or 802.1w) switch or group of switches. (Where multiple paths exist between an MST region and a legacy switch, expect the CST (Common Spanning Tree) to block all but one such path.)
- Determine the root bridge and root port for each instance.
- Determine the designated bridge and designated port for each LAN segment.
- Determine which VLANs to assign to each instance and use port trunks with 802.1Q VLAN tagging where separate links for separate VLANs would result in a blocked link preventing communication between nodes on the same VLAN. (See **MSTP operation with 802.1Q VLANs** .)
- Identify the edge ports connected to end nodes and enable the `admin-edge-port` setting for these ports. Leave the admin-edge-port setting disabled for ports connected to another switch, a bridge, or a half-duplex repeater.
- When there is a DHCP assigned IP address on a switch with only one active port and Spanning Tree is enabled, the port is blocked (initial state of STP convergence). The VLAN goes down and the DHCP assigned IP address is removed from the VLAN. After STP converges, the port becomes forwarding and the DHCP IP address is restored.

Configuring MSTP at a glance

The general steps for configuring MSTP via the CLI are:

Procedure

1. Configure MSTP global parameters. This involves:

   a. Selecting MSTP as the spanning tree mode:

   ```
   spanning-tree mode mstp
   ```

   b. Clearing spanning tree debug counters:

   ```
   spanning-tree clear-debug-counters
   ```

   c. Specifying required parameters for MST region identity:

   ```
   Region Name:spanning-tree config-name Region Revision Number:spanning-tree
   config-revision
   ```

   d. Optionally, specifying MSTP parameter changes for region settings:

HPE recommends that you leave these parameters at their default settings for most networks. See the Caution below.

- The maximum number of hops before the MSTP BPDU (Bridge Protocol Data Unit) is discarded: `spanning-tree max-hops` (default: 20)
- Force-Version operation: `spanning-tree force-version`
- Forward Delay: `spanning-tree forward-delay`
- Hello Time (if it is the root device): `spanning-tree hello-time`
- Maximum age to allow for STP packets before discarding: `spanning-tree maximum-age`
- Device spanning tree priority. Specifies the priority value used along with the switch MAC address to determine which device is root. The lower a priority value, the higher the priority: `spanning-tree priority`

e. Enabling SNMP traps:

```
no spanning-tree trap [ errant-bpdu | loop-guard | new-root | root-guard ]
```

CAUTION: When you enable MSTP on the switch, the default MSTP spanning tree configuration settings comply with the values recommended in the IEEE 802.1s Multiple Spanning Tree Protocol (MSTP) standard. Inappropriate changes to these settings can result in severely degraded network performance. For this reason, HPE strongly recommends that changing these default settings be reserved only for experienced network administrators who have a strong understanding of the IEEE 802.1D/w/s standards and operation.

2. Configure per port parameters. HPE recommends that you use the default settings for these parameters and apply changes on a per-port basis only where a nondefault setting is clearly indicated by the circumstances of individual links. Other features you might consider include BPDU Filtering or BPDU Protection—these provide additional per-port control over spanning tree operations and security on the switch.

3. Configure MST instances. Configure one instance for each VLAN group that you want to operate as an active topology within the region to which the switch belongs. When you create the instance, you must include a minimum of one VID. You can add more VIDs later if desired. Use the following command:

```
spanning-tree instance n vlan vid
```

To move a VLAN from one instance to another, first use `no spanning-tree instance` *n* `vlan` *vid* to remove the mapping from the VLAN from the current instance, then add the VLAN to the other instance. (While the VLAN mapping is removed from an MSTI, it is associated with the region's IST instance.)

4. Configure the priority for each instance with the following command: `spanning-tree instance` *n* `priority` *n*

5. Configure MST instance port parameters. HPE recommends that you apply changes on a per-port basis only where a nondefault setting is clearly indicated by the circumstances of individual links. For example, you might want to set the path cost value for the *ist* or for the ports used by a specific MST instance. Use the following command:`spanning-tree instance <ist> | 1..64 <port-list> path-cost [ auto | 1..200000000 ]` Alternatively, leaving this setting at the default (auto) allows the switch to calculate the path-cost from the link speed.

6. Enable spanning tree operation on the switch with the `spanning-tree` command.

Configuring MSTP operation mode and global settings

The commands in this section apply at the switch (global) level.

Selecting MSTP as the spanning tree mode

Syntax:

```
spanning-tree mode mstp
```

Specifies that spanning tree will run in MSTP mode.

Clearing spanning tree debug counters

Syntax:

```
spanning-tree clear-debug-counters
```

Clears spanning tree debug counters.

Resetting the configuration name of the MST region in which a switch resides

Syntax:

```
spanning-tree config-name ascii-string
```

```
no spanning-tree config-name ascii-string
```

Resets the configuration name of the MST region in which the switch resides. This name can include up to 32 nonblank characters and is case-sensitive. On all switches within a given MST region, the configuration names must be identical. For more than one MSTP switch in the same MST region, the identical region name must be configured on all such switches. If the default configuration name is retained on a switch, it cannot exist in the same MST region with another switch.

The default name is a text string using the hexadecimal representation of the switch's MAC address.

The `no` form of the command overwrites the currently configured name with the default name.

NOTE: This option is available only when the switch is configured for MSTP operation. There is no defined limit on the number of regions you can configure.

Designating the revision number of the MST region for a switch

Syntax:

```
spanning-tree config-revision revision-number
```

Configures the revision number designated for the MST region in which you want the switch to reside. This setting must be the same for all switches residing in the same region. Use this setting to differentiate between region configurations in situations such as the following:

- Changing configuration settings within a region where you want to track the configuration versions you use
- Creating a new region from a subset of switches in a current region and want to maintain the same region name.
- Using the `pending` option to maintain two different configuration options for the same physical region.

This setting must be the same for all MSTP switches in the same MST region.

Range: 0 - 65535

Default: 0

NOTE:

This option is available only when the switch is configured for MSTP operation.

Setting the spanning tree compatibility mode

Syntax:

```
spanning-tree force-version [ stp-compatible | rstp-operation | mstp-operation ]
```

Sets the spanning tree compatibility mode. This command forces the switch to emulate behavior of earlier versions of spanning tree protocol, or return to MSTP behavior. The command is useful in test or debug applications, and removes the need to reconfigure the switch for temporary changes in spanning tree operation.

`stp-compatible`

The switch applies 802.1D STP operation on all ports.

`rstp-operation`

The switch applies 802.1w operation on all ports except those ports where it detects a system using 802.1D Spanning Tree. RSTP is Rapid Spanning Tree Protocol.

`mstp-operation`

The switch applies 802.1s MSTP operation on all ports where compatibility with 802.1D or 802.1w spanning tree protocols is not required.

NOTE:

Even when `mstp-operation` is selected, if the switch detects an 802.1D BPDU or an 802.1w BPDU on a port, it communicates with the device linked to that port using STP or RSTP BPDU packets. Also, if errors are encountered as described in **Configuring MSTP at a glance** on page 114, setting `force-version` to `stp-compatible` forces the MSTP switch to communicate out all ports using operations that are compatible with IEEE 802.1D STP.

NOTE:

When using MSTP rapid state transitions

Under some circumstances the rapid state transitions employed by MSTP can increase the rates of frame duplication and incorrect ordering in the switched LAN. To allow MSTP switches to support applications and protocols that may be sensitive to frame duplication and incorrect ordering, setting the Force Protocol Version (`force-version`) parameter to `stp-compatible` allows MSTP to operate with rapid transitions disabled. The value of this parameter applies to all ports on the switch.

Setting the time interval between listening, learning, and forwarding states

Syntax:

```
spanning-tree forward-delay
```

Sets the time the switch waits between transitions from listening to learning and from learning to forwarding states.

Range: 4 - 30

Default: 15 seconds

Setting spanning tree to operate in 802.1D legacy mode

Syntax:

```
spanning-tree legacy-mode
```

```
no spanning-tree legacy-mode
```

Forces spanning tree to operate in legacy (802.1D) mode.

Default: MSTP-operation.

The `no` form of this command returns the switch to the default 802.1s native mode (MSTP-operation.)

Setting spanning tree to operate with 802.1D legacy path cost values

Syntax:

```
spanning-tree legacy-path-cost
```

```
no spanning-tree legacy-path-cost
```

Forces spanning tree to operate with legacy (802.1D) path cost values.

Default: 802.1t.

The `no` form of the command returns the switch to the default 802.1t (not legacy) path cost values.

Specifying the time interval between BPDU transmissions

Syntax:

```
spanning-tree hello-time 1..10
```

If MSTP is running and the switch is operating as the CIST (Common and Internal Spanning Tree) root for your network, this command specifies the time in seconds between transmissions of BPDUs for all ports on the switch configured with the Global option (the default). This parameter applies in MSTP, RSTP, and STP modes.

During MSTP operation, you can override this global setting on a per-port basis with this command: `spanning-tree` *`port-list`* `hello-time` *`1..10`* .

Default: 2 seconds.

Setting the hop limit for BPDUs

Syntax:

```
spanning-tree max-hops hop-count
```

Resets the number of hops allowed for BPDUs in an MST region. When an MSTP switch receives a BPDU, it decrements the hop-count setting the BPDU carries. If the hop-count reaches zero, the receiving switch drops the BPDU.

The switch does not change the message-age and maximum-age data carried in the BPDU as it moves through the MST region and is propagated to other regions.

Range: 1 - 40

Default: 20

Setting the maximum age of received STP information

Syntax:

```
spanning-tree maximum age
```

Sets the maximum age time for received STP information before it is discarded.

Default: 20 seconds

Manipulating the pending MSTP configuration

Syntax:

```
spanning-tree pending [apply | config-name | config-revision | instance | reset]
```

Manipulates the pending MSTP configuration. The command is useful in test or debug applications, and enables rapid reconfiguration of the switch for changes in spanning tree operation.

`apply`

Applies pending MSTP configuration (swaps active and pending configurations).

`config-name`

Sets the pending MST region configuration name. Default is the switch's MAC address.

`config-revision`

Sets the pending MST region configuration revision number. Default is 0.

`instance`

Change pending MST instance configuration.

`reset`

Copies the active configuration to pending.

Setting the bridge priority for a region and determining the root switch

Syntax:

```
spanning-tree priority priority-multiplier
```

Every switch running an instance of MSTP has a Bridge Identifier, which is a unique identifier that helps distinguish this switch from all others. The switch with the lowest Bridge Identifier is elected as the root for the tree.

The Bridge Identifier is composed of a configurable priority component (2 bytes) and the bridge's MAC address (6 bytes). You can change the priority component provides flexibility in determining which switch will be the root for the tree, regardless of its MAC address.

This command sets the switch (bridge) priority for the designated region in which the switch resides. The switch compares this priority with the priorities of other switches in the same region to determine the root switch for the region. The lower the priority value, the higher the priority. If there is only one switch in the region, then that switch is the root switch for the region. The root bridge in a region provides the path to connected regions for the traffic in VLANs assigned to the region's IST instance. Traffic in VLANs assigned to a numbered STP instance in a given region moves to other regions through the root switch for that instance.

The priority range for an MSTP switch is 0-61440. However, this command specifies the priority as a multiplier (0 - 15) of 4096. That is, when you specify a priority multiplier value of 0 - 15, the actual priority assigned to the switch is: (priority-multiplier) x 4096

For example, with 2 as the priority-multiplier on a given MSTP switch, the Switch Priority setting is 8,192.

NOTE:

If multiple switches in the same MST region have the same priority setting, then the switch with the lowest MAC address becomes the root switch for that region.

Enabling SNMP traps

Syntax:

```
spanning-tree trap {errant-bpdu | loop-guard | new-root | root-guard}
```

```
no spanning-tree trap {errant-bpdu | loop-guard | new-root | root-guard}
```

Enables or disables SNMP traps for errant-BPDU, loop guard, new root, and root guard event notifications. This command is designed to be used in conjunction with the `spanning-tree bpdu-filter` command and the `bpdu-protection` command.

Parameters

`errant-bpdu`

Enables SNMP notification when an errant BPDU is received. Designed for use with BPDU filtering.

`loop-guard`

Enables SNMP notification when a loop guard inconsistency is detected. Designed for use with the Loop Guard option.

`new-root`

Enables SNMP notification when a new root is elected on any VLAN configured for MSTP on the switch.

`root-guard`

Enables SNMP notification when a root guard inconsistency is detected.

Default for all of the above options: Disabled

The `no` form of the command disables traps on the switch.

Configuring MSTP per-port parameters

In an MSTP topology, per-port parameters are set in the global configuration context. In most cases, HPE recommends that you use the default settings for these parameters and apply changes on a per-port basis only where a nondefault setting is clearly indicated by the circumstances of individual links. Some port parameters (such as `admin-edge-port`) affect all MSTI instances that consist of VLANs configured on the port. Other port parameters (such as `path-cost`) affect only the specified MST.

Enabling immediate transition to forwarding on end nodes

Syntax:

```
spanning-tree port-list admin-edge-port
```

```
no spanning-tree port-list admin-edge-port
```

Enables `admin-edge-port` on ports connected to end nodes. During spanning tree establishment, ports with `admin-edge-port` enabled transition immediately to the forwarding state. If a bridge or switch is detected on the segment, the port automatically operates as non-edge, not enabled.

Default: Disabled

If admin-edge-port is disabled on a port and auto-edge-port has not been disabled, the auto-edge-port setting controls the behavior of the port.

The no form of this command disables edge port operation on the specified ports.

Identifying edge ports automatically

Syntax:

```
spanning-tree port-list auto-edge-port
```

```
no spanning-tree port-list auto-edge-port
```

Enables automatic identification of edge ports for faster convergence. When enabled, the port looks for BPDUs for the first 3 seconds. If there are none, the port is classified as an edge port and immediately starts forwarding packets. If BPDUs are seen on the port, the port is classified as a non-edge port and normal STP operation commences on that port.

If admin-edge-port is enabled for a port, the setting for auto-edge-port is ignored whether set to yes or no.

If admin-edge-port is set to no, and auto-edge-port has not been disabled (set to no), then the auto-edge-port setting controls the behavior of the port.

CAUTION: Requires thorough knowledge of IEEE 802.1D/w/s standards and operation.

Spanning tree interprets a switch mesh as a single link. Because the switch automatically gives faster links a higher priority, the default MSTP parameter settings are usually adequate for spanning tree operation. Because incorrect MSTP settings can adversely affect network performance, do not change the MSTP settings from their default values unless you have a strong understanding of how spanning tree operates.

Default: Enabled

The no form of this command disables auto-edge-port operation on the specified ports.

Specifying the interval between BPDU transmissions

Syntax:

```
spanning-tree port-list hello-time [global | 1 - 10]
```

When the switch is the CIST root, this parameter specifies the interval (in seconds) between periodic BPDU transmissions by the designated ports. This interval also applies to all ports in all switches downstream from each port in the *port-list*.

A setting of global indicates that the ports in *port-list* on the CIST root are using the value set by the global spanning tree hello-time value.

When a given switch X is not the CIST root, the per-port hello-time for all active ports on switch X is propagated from the CIST root, and is the same as the hello-time in use on the CIST root port in the currently active path from switch X to the CIST root. When switch X is not the CIST root, then the upstream CIST root's port hello-time setting overrides the hello-time setting configured on switch X.

Default Per-Port setting: Use Global.

Default Global Hello-Time: 2.

Forcing a port to send RST/MST BPDUs

Syntax:

```
spanning-tree port-list mcheck
```

Forces a port to send RST/MST BPDUs for 3 seconds. This tests whether all STP bridges on the attached LAN have been removed and the port can migrate to native MSTP mode and use RST/MST BPDUs for transmission.

Determining which ports are forwarding ports by assigning port cost

Syntax:

```
spanning-tree port-list path-cost [auto | 1..200000000]
```

Assigns an individual port cost that the switch uses to determine which ports are forwarding ports in a given spanning tree. In the default configuration (auto) the switch determines a port's path cost by the port's type:

10 Mbps

`2000000`

100 Mbps

`200000`

1 Gbps

`20000`

Default: Auto

Informing the switch of the device type to which a port connects

Syntax:

```
spanning-tree port-list point-to-point-mac [true | false | auto]
```

Informs the switch of the type of device to which a specific port connects.

Parameters

`true`

(Default) Indicates a point-to-point link to a device such as a switch, bridge, or end-node.

`false`

Indicates a connection to a half-duplex repeater (which is a shared LAN segment).

`auto`

Causes the switch to set Force-False on the port if it is not running at full duplex.

Determining which port to use for forwarding

Syntax:

```
spanning-tree port-list priority priority-multiplier
```

MSTP uses this parameter to determine the port to use for forwarding. The port with the lowest priority number has the highest priority for use.

The range is 0 to 240, and is configured by specifying a multiplier from 0 - 15. When you specify a priority multiplier of 0 - 15, the actual priority assigned to the switch is:

(priority-multiplier) x 16

If you configure `2` as the priority multiplier on a given port, the actual Priority setting is 32. After specifying the port priority multiplier, the switch displays the actual port priority (and not the multiplier) in the `show spanning-tree` or `show spanning-tree port-list` displays.

You can view the actual multiplier setting for ports by executing `show running` and looking for an entry in this format:

```
spanning-treeport-list priority priority-multiplier
```

For example, configuring port A2 with a priority multiplier of `3` results in the following line in the `show running` output:

```
spanning-tree A2 priority 3
```

Denying a port the role of root port

Syntax:

```
spanning-treeport-list root-guard
```

When a port is enabled as `root-guard`, it cannot be selected as the root port even if it receives superior STP BPDUs. The port is assigned an "alternate" port role and enters a blocking state if it receives superior STP BPDUs.

A superior BPDU contains both "better" information on the root bridge and path cost to the root bridge, which would normally replace the current root bridge selection.

The superior BPDUs received on a port enabled as `root-guard` are ignored. All other BPDUs are accepted and the external devices may belong to the spanning tree as long as they do not claim to be the Root device.

Use this command on MSTP switch ports that are connected to devices located in other administrative network domains to:

- Ensure the stability of the core MSTP network topology so that undesired or damaging influences external to the network do not enter.
- Protect the configuration of the CIST root bridge that serves as the common root for the entire network.

Default: Disabled

Denying a port propagation change information

Syntax:

```
spanning-tree port-list tcn-guard
```

When enabled for a port, this causes the port to stop propagating received topology change notifications and topology changes to other ports.

Default: Disabled

Configure MST instance ports parameters

Syntax

```
spanning-tree instance 1-64 ethernet PORT-LIST
```

Description

Configure MST instance ports parameters.

Best practices

Follow the PORT-LIST with a '?' to get the list of all possible options.

Create a new instance or map VLAN(s) to an existing one

Syntax

```
spanning-tree instance ist | 1-64 vlan VLAN-ID
```

Description

Used to create a new instance or map VLAN(s) to an existing one. Each instance must have at least one VLAN mapped to it. The VLANs unmapped from other instances are automatically mapped to the IST instance. Only IST VLANs can be directly mapped to other instances. When VLANs are mapped to an instance, they are automatically unmapped from the instance they were mapped to before. Any MSTP instance can have all the VLANs configured in the switch.

Enable event logging

Syntax

```
no spanning-tree log state-transitions instance 1-64 | ist vlan
```

Description

By default port state change for IST is added in log.

Deleting an instance

Syntax

```
no spanning-tree instance <1-64>
```

Description

Deletes an instance. The IST instance cannot be deleted.

Configure an existent instance

Syntax

```
no spanning-tree instance <ist|1-64>
```

Description

Used to configure an existent instance.

Best Practices

Follow the syntax with a '?' to get a complete list of all the configurable parameters and sub-commands.

MSTP Config example

MSTP Config

```
VLAN 109
  ip addr 10.0.109.10/24
  tag 12
  exit

VLAN 110
  ip addr 10.0.110.10/24
  tag 12
  exit

Spanning-tree
Spanning-tree mode mstp
Spanning-tree config-name "MSTPRegion1"
Spanning-tree config-revision 1
Spanning-tree instance 1 VLAN 109
Spanning-tree instance 1 priority 4
Spanning-tree instance 2 VLAN 110
```

Downgrading to lower version build

The downgrade to lower version build will result in “stuck in boot” if more than 16 instances are created in the DUT.

Operating notes for the VLAN configuration enhancement

- Configuring MSTP on the switch automatically configures the Internal Spanning Tree (IST) instance and places all statically and dynamically configured VLANs on the switch into the IST instance. The spanning tree instance vlan command creates a new MST instance and moves the VLANs you specify from the IST to the MSTI.You must map a least one VLAN ID to an MSTI when you create it. You cannot map a VLAN ID to more than one instance. You can create up to 16 MSTIs in a region.
- The `no` form of the spanning tree instance vlan command removes one or more VLANs from the specified MSTI. If no VLANs are specified, the `no` form of the command deletes the specified MSTI.When you remove a VLAN from an MSTI, the VLAN returns to the IST instance, where it can remain or be reassigned to another MSTI configured in the region.
- If you enter the spanning tree instance vlan command before a static or dynamic VLAN is configured on the switch to preconfigure VLAN ID-to-MSTI mappings, no error message is displayed. Later, each newly configured VLAN that has already been associated with an MSTI is automatically assigned to the MSTI.This new default behavior differs from automatically including configured (static and dynamic) VLANs in the IST instance and requiring you to manually assign individual static VLANs to an MSTI.
- The valid VLAN IDs that you can map to a specified MSTI are from 1 to 4094. The VLAN ID-to-MSTI mapping does not require a VLAN to be already configured on the switch. The MSTP VLAN enhancement allows you to preconfigure MSTP topologies before the VLAN IDs associated with each instance exist on a switch.

- When you use preconfigured VLAN ID-to-MSTI topologies, ensure that MSTP switches remain in the same region by mapping all VLAN IDs used in the region to the same MSTIs on each regional switch.
- When you upgrade switch software to release K.13.XX and later, the existing MSTP topology configuration is automatically saved. All existing VLAN ID-to-MSTI assignments are maintained on a switch for uninterrupted MSTP network operation.

Configuring MST instance parameters

When you enable MSTP on the switch, a spanning tree instance is enabled automatically. The switch supports up to 16 configurable MST instances for each VLAN group that you want to operate as an active topology within the region to which the switch belongs. When creating an instance, you must include a minimum of one VID. You can add more VIDs later if desired.

Syntax:

```
spanning-tree instance 1..16 vlan vid [vid..vid]
```

```
no spanning-tree instance 1..16 vlan vid [vid..vid]
```

Configuring MSTP on the switch automatically configures the IST instance and places all statically and dynamically configured VLANs on the switch into the IST instance. This command creates a new MST instance (MSTI) and moves the VLANs you specify from the IST to the MSTI.

You must map at least one VLAN to an MSTI when you create it. You cannot map a VLAN ID to more than one instance. You can create up to 16 MSTIs in a region.

The `no` form of the command removes one or more VLANs from the specified MSTI. If no VLANs are specified, the `no` form of the command deletes the specified MSTI.

When you remove a VLAN from an MSTI, the VLAN returns to the IST instance, where it can remain or be re-assigned to another MSTI configured in the region.

NOTE: Starting in software release 13.x.x, you can enter the `spanning-tree instance vlan` command before a static or dynamic VLAN is configured on the switch to preconfigure VLAN ID-to-MSTI mappings. No error message is displayed. Later, each newly configured VLAN that has already been associated with an MSTI is automatically assigned to the MSTI.

This new default behavior differs from automatically including configured (static and dynamic) VLANs in the IST instance and requiring the manual assigning of individual static VLANs to an MSTI.

NOTE: The valid VLAN IDs that you can map to a specified MSTI are from 1 to 4094. The VLAN ID-to-MSTI mapping does not require a VLAN to be already configured on the switch. The MSTP VLAN enhancement allows you to preconfigure MSTP topologies before the VLAN IDs associated with each instance exist on a switch.

When using preconfigured VLAN ID-to-MSTI topologies, be sure that MSTP switches remain in the same region by mapping all VLAN IDs used in the region to the same MSTIs on each regional switch.

When upgrading switch software to release 13.x.x and later, the existing MSTP topology configuration is automatically saved. All existing VLAN ID-to-MSTI assignments are maintained on a switch for uninterrupted MSTP network operation.

Setting the bridge priority for an instance

Syntax:

```
spanning-tree instance 1..16 priority priority-multiplier
```

Sets the switch (bridge) priority for the designated instance. This priority is compared with the priorities of other switches in the same instance to determine the root switch. The lower the priority value, the higher the priority. If there is only one switch in the instance, then that switch is the root switch for the instance. The IST regional root bridge provides the path to instances in other regions that share one or more of the same VLANs.

The priority range for an MSTP switch is 0 - 61440. However, this command specifies the priority as a multiplier (0 - 15) of 4096. When a priority multiplier value is set from 0 - 15, the actual priority assigned to the switch for the specified MST instance is: (priority-multiplier) x 4096

For example, if you configure 5 as the priority-multiplier for MST Instance 1 on a given MSTP switch, the Switch Priority setting is 20,480 for that instance in that switch.

NOTE:

If multiple switches in the same MST instance have the same priority setting, the switch with the lowest MAC address becomes the root switch for that instance.

Assigning a port cost for an MST instance

Syntax:

```
spanning-tree instance ist | 1..16 port-list path-cost [auto | 1..200000000]
```

Assigns an individual port cost for the IST or for the specified MST instance.

For a given port, the path cost setting can be different for different MST instances to which the port may belong. The switch uses the path cost to determine which ports are the forwarding ports in the instance; that is, which links to use for the active topology of the instance and which ports to block.

The settings are either `auto` or in a range from 1 to 200,000,000. With the `auto` setting, the switch calculates the path cost from the link speed:

10 Mbps

`2000000`

100 Mbps

`200000`

1 Gbps

`20000`

Default

`Auto`

Setting the priority for a port in a specified MST instance

Syntax:

```
spanning-tree instance 1..16 port-list priority priority-multiplier
```

Sets the priority for the specified ports in the specified MST instance.

For a given port, the priority setting can be different for different MST instances to which the port may belong. The priority range for a port in a given MST instance is 0 - 255. However, this command specifies the priority as a multiplier (0 - 15) of 16. When you specify a priority multiplier of 0 - 15, the actual priority assigned to the switch is: (priority-multiplier) x 16

Setting priority for a port in a specified MST instance

If you configure 2 as the priority multiplier on a given port in an MST instance, then the actual Priority setting is 32x. After you specify the port priority multiplier in an instance, the switch displays the actual port priority and not the multiplier in the `show spanning-tree instance 1..16` or `show spanning-tree port-list instance 1..16` displays.

You can view the actual multiplier setting for ports in the specified instance by executing `show running` and looking for an entry in the following format:

```
spanning-tree instance 1..15 port-list priority priority-multiplier
```

For example, configuring port A2 with a priority multiplier of `3` in instance 1, results in this line in the `show running` output:

```
spanning-tree instance 1 A2 priority 3
```

Setting the priority for specified ports for the IST

Syntax:

```
spanning-tree port-list priority priority-multiplier
```

Sets the priority for the specified ports for the IST (Instance 0) of the region in which the switch resides.

The priority component of the port's Port Identifier is set. The Port Identifier is a unique identifier that helps distinguish this switch's ports from all others. It consists of the priority value with the port number extension—PRIORITY:PORT_NUMBER. A port with a lower value of Port Identifier is more likely to be included in the active topology.

This priority is compared with the priorities of other ports in the IST to determine which port is the root port for the IST instance. The lower the priority value, the higher the priority. The IST root port (or trunk) in a region provides the path to connected regions for the traffic in VLANs assigned to the region's IST instance.

The priority range for a port in a given MST instance is 0 - 240. However, this command specifies the priority as a multiplier (0 - 15) of 16. When you specify a priority multiplier of 0 - 15, the actual priority assigned to the switch is: (priority-multiplier) x 16

Setting priority for specified ports for an IST

Configuring 5 as the priority multiplier on a given port in the IST instance for a region creates an actual priority setting of 80. After specifying the port priority multiplier for the IST instance, the switch displays the actual port priority, not the multiplier, in the `show spanning-tree instance ist` or `show spanning-tree port-list instance ist` displays. You can view the actual multiplier setting for ports in the IST instance by executing `show running` and looking for an entry in this format:

```
spanning-tree port-list priority priority-multiplier
```

So configuring port A2 with a priority multiplier of `2` in the IST instance, results in this line in the `show running` output:

```
spanning-tree A2 priority 2
```

Enabling or disabling spanning tree operation

This command enables or disables spanning tree operation for any spanning tree protocol enabled on the switch. Before using to enable spanning tree, be sure that the right version is active on the switch.

Syntax:

```
no spanning-tree
```

Enables or disables spanning tree. Enabling spanning tree with MSTP configured, implements MSTP for all physical ports on the switch according to the VLAN groupings for the IST instance and any other configured instances.

Disabling MSTP removes protection against redundant loops that can significantly slow or halt a network.

This command simply turns spanning tree on or off. It does not change the existing spanning tree configuration.

NOTE: The convergence time for implementing MSTP changes can be disruptive to your network. To minimize such disruption, consider using the `spanning-tree pending` command.

Enabling an entire MST region at once or exchanging one region configuration for another

This operation exchanges the currently active MSTP configuration with the currently pending MSTP configuration, making it possible to implement a new MSTP configuration with minimal network disruption or to exchange MSTP configurations for testing or troubleshooting purposes.

When configuring or reconfiguring MSTP, the switch recalculates the corresponding network paths. This can have a ripple effect throughout your network as adjacent MSTP switches recalculate network paths to support the configuration changes invoked in a single switch. Although MSTP employs rapid spanning tree operation, the convergence time for implementing MSTP changes can be disruptive to your network. However, by using the `spanning-tree pending` feature, you can set up an MSTP on the switch and then invoke all instances of the new configuration at the same time, instead of one at a time.

Syntax:

```
no spanning-tree pending [apply | config-name | config-revision | instance | reset]
```

Exchanges the currently active MSTP configuration with the current pending MSTP configuration. Options are as follows:

`apply`

Exchanges the currently active MSTP configuration with the pending MSTP configuration.

`config-name`

Specifies the pending MST region name. Must be the same for all MSTP switches in the region. (Default: The switch's MAC address.)

`config-revision`

Specifies the pending MST region configuration revision number. Must be the same for all MSTP switches in the region. (Default: 0).

`instance`

Creates the pending instance and assigns one or more VLANs to the instance.

`reset`

Copies the switch's currently active MSTP configuration to the pending configuration. This is useful when you want to experiment with the current MSTP configuration while maintaining an unchanged version.

Creating a pending MSTP configuration

To create a pending MSTP configuration and exchange it with the active MSTP configuration:

Procedure

1. Configure the VLANs to include in any instances in the new region. When you execute the `pending` command, all VLANs configured on the switch will be assigned to a single pending IST instance unless assigned to other, pending MST instances. The `pending` command creates the region's IST instance automatically.
2. Configure MSTP as the spanning tree protocol, then execute `write mem` and reboot. The pending option is available only with MSTP enabled.
3. Configure the pending region *config-name* to assign to the switch.
4. Configure the pending *config-revision* number for the region name.
5. If you want an MST instance other than the IST instance, configure the instance number and assign the appropriate VLANs (VIDs) using the `pending instance`*1..16*`vlan [vid |` *vid-range*`] command.`
6. Repeat step 5 for each additional MST instance necessary.
7. To review your pending configuration, use the `show spanning-tree pending` command.
8. To exchange the currently active MSTP configuration with the pending MSTP configuration, use the `spanning-tree pending apply` command.

Viewing MSTP statistics

NOTE:

SNMP MIB Support for MSTP is a superset of the STP/802.1D and RSTP/802.1w protocols and uses the MIB objects defined for these two protocols.

Viewing global MSTP status

The following commands display the MSTP statistics for the connections between MST regions in a network.

Syntax:

```
show spanning-tree
```

Displays the switch's global and regional spanning tree status, plus the per-port spanning tree operation at the regional level. Values for the following parameters appear only for ports connected to active devices: `Designated Bridge, Hello Time, PtP, and Edge.`

Syntax:

`show spanning-tree` *port-list*

Displays the spanning tree status for the designated ports. You can list data for a series of ports and port trunks by specifying the first and last port or trunk of any consecutive series of ports and trunks. For example, to display data for port A20-A24 and trk1, you would use this command: `show spanning-tree a20-a42,trk1`

Viewing a common spanning tree status

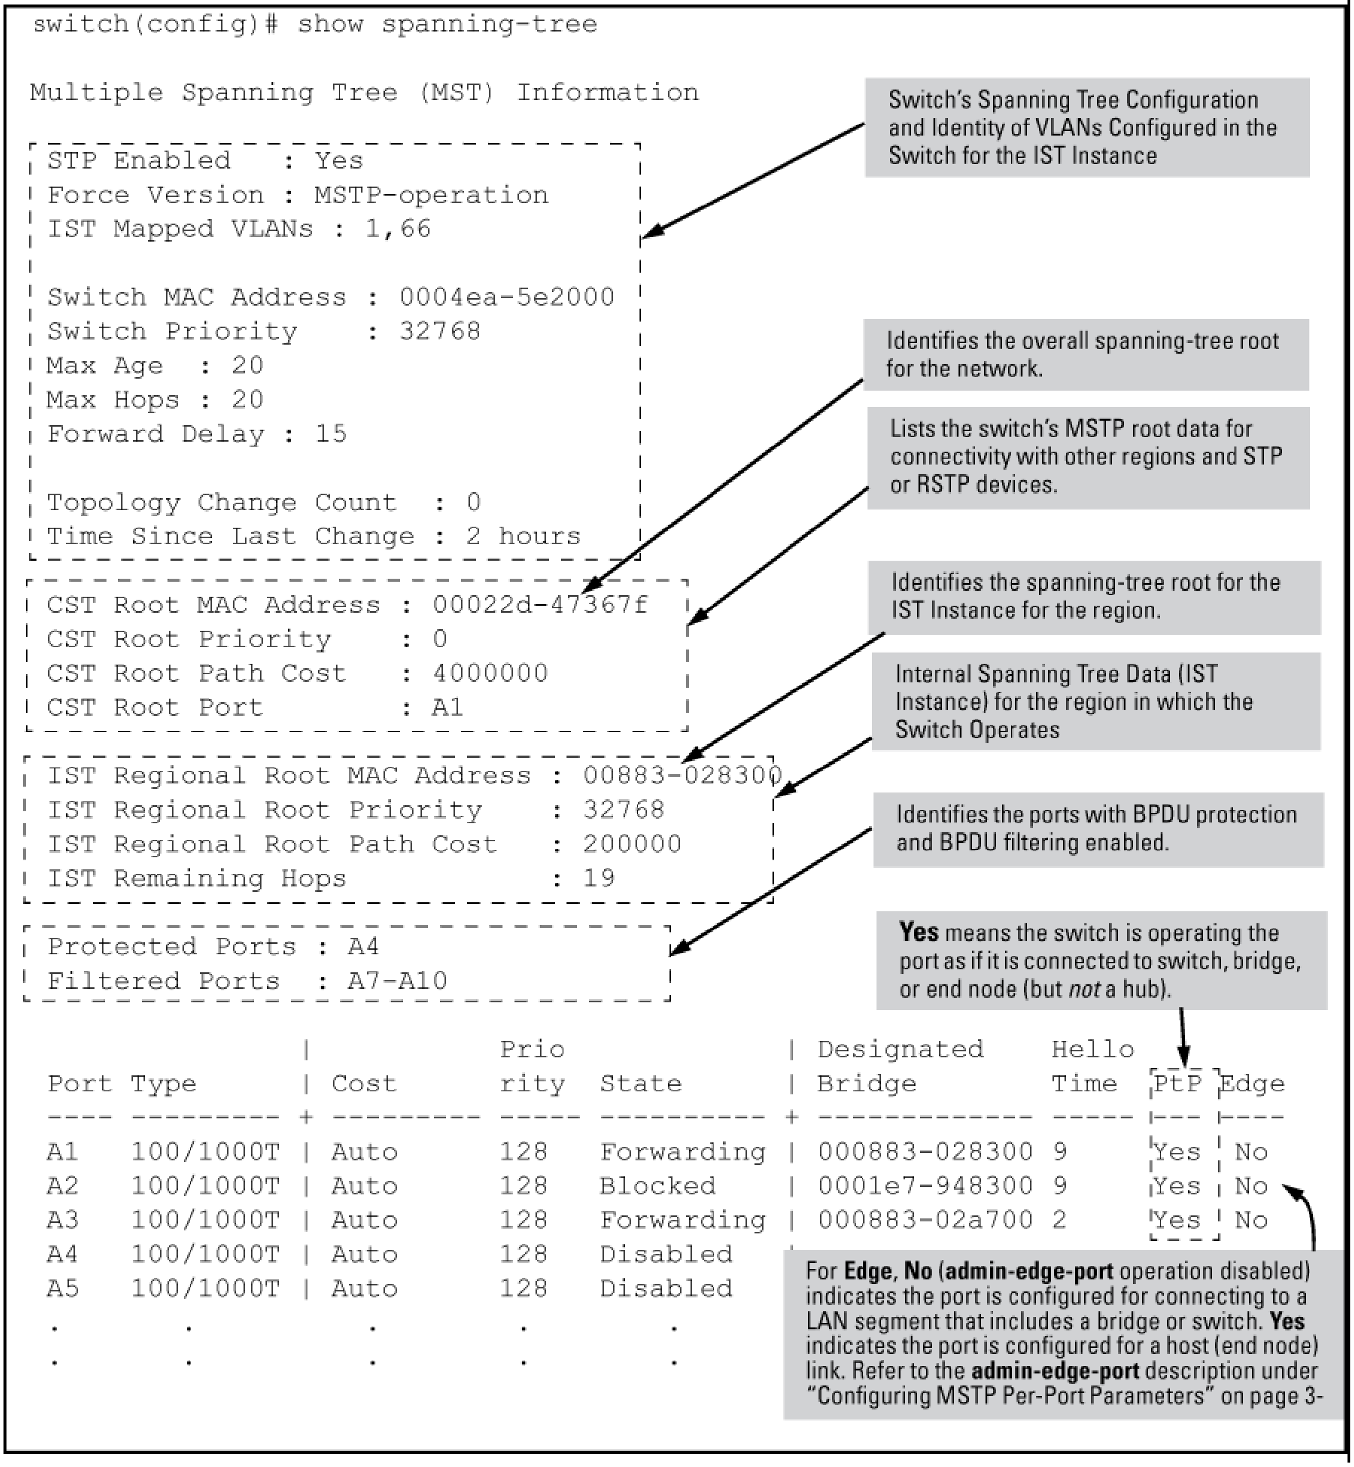

```
switch(config)# show spanning-tree

Multiple Spanning Tree (MST) Information

 STP Enabled    : Yes
 Force Version : MSTP-operation
 IST Mapped VLANs : 1,66

 Switch MAC Address : 0004ea-5e2000
 Switch Priority     : 32768
 Max Age  : 20
 Max Hops : 20
 Forward Delay : 15

 Topology Change Count  : 0
 Time Since Last Change : 2 hours

 CST Root MAC Address : 00022d-47367f
 CST Root Priority    : 0
 CST Root Path Cost   : 4000000
 CST Root Port        : A1

 IST Regional Root MAC Address : 00883-028300
 IST Regional Root Priority    : 32768
 IST Regional Root Path Cost   : 200000
 IST Remaining Hops            : 19

 Protected Ports : A4
 Filtered Ports  : A7-A10

                    |           Prio                  | Designated    Hello
 Port Type          | Cost      rity  State           | Bridge        Time  PtP  Edge
 ---- --------- + --------- ----  ----------  + ------------- ----- ---  ----
 A1   100/1000T | Auto      128   Forwarding  | 000883-028300 9     Yes  No
 A2   100/1000T | Auto      128   Blocked     | 0001e7-948300 9     Yes  No
 A3   100/1000T | Auto      128   Forwarding  | 000883-02a700 2     Yes  No
 A4   100/1000T | Auto      128   Disabled    |
 A5   100/1000T | Auto      128   Disabled
 .        .          .       .         .
 .        .          .       .         .
```

Switch's Spanning Tree Configuration and Identity of VLANs Configured in the Switch for the IST Instance

Identifies the overall spanning-tree root for the network.

Lists the switch's MSTP root data for connectivity with other regions and STP or RSTP devices.

Identifies the spanning-tree root for the IST Instance for the region.

Internal Spanning Tree Data (IST Instance) for the region in which the Switch Operates

Identifies the ports with BPDU protection and BPDU filtering enabled.

Yes means the switch is operating the port as if it is connected to switch, bridge, or end node (but *not* a hub).

For **Edge**, **No** (**admin-edge-port** operation disabled) indicates the port is configured for connecting to a LAN segment that includes a bridge or switch. **Yes** indicates the port is configured for a host (end node) link. Refer to the **admin-edge-port** description under "Configuring MSTP Per-Port Parameters" on page 3-

Viewing detailed port information

The following commands display the MSTP statistics for the connections between MST regions in a network.

Syntax:

```
show spanning-tree detail
```

Displays additional parameters concerning the CST ports.

Syntax:

```
show spanning-tree port-list detail
```

Displays detailed spanning tree status for the designated ports.

Viewing port information

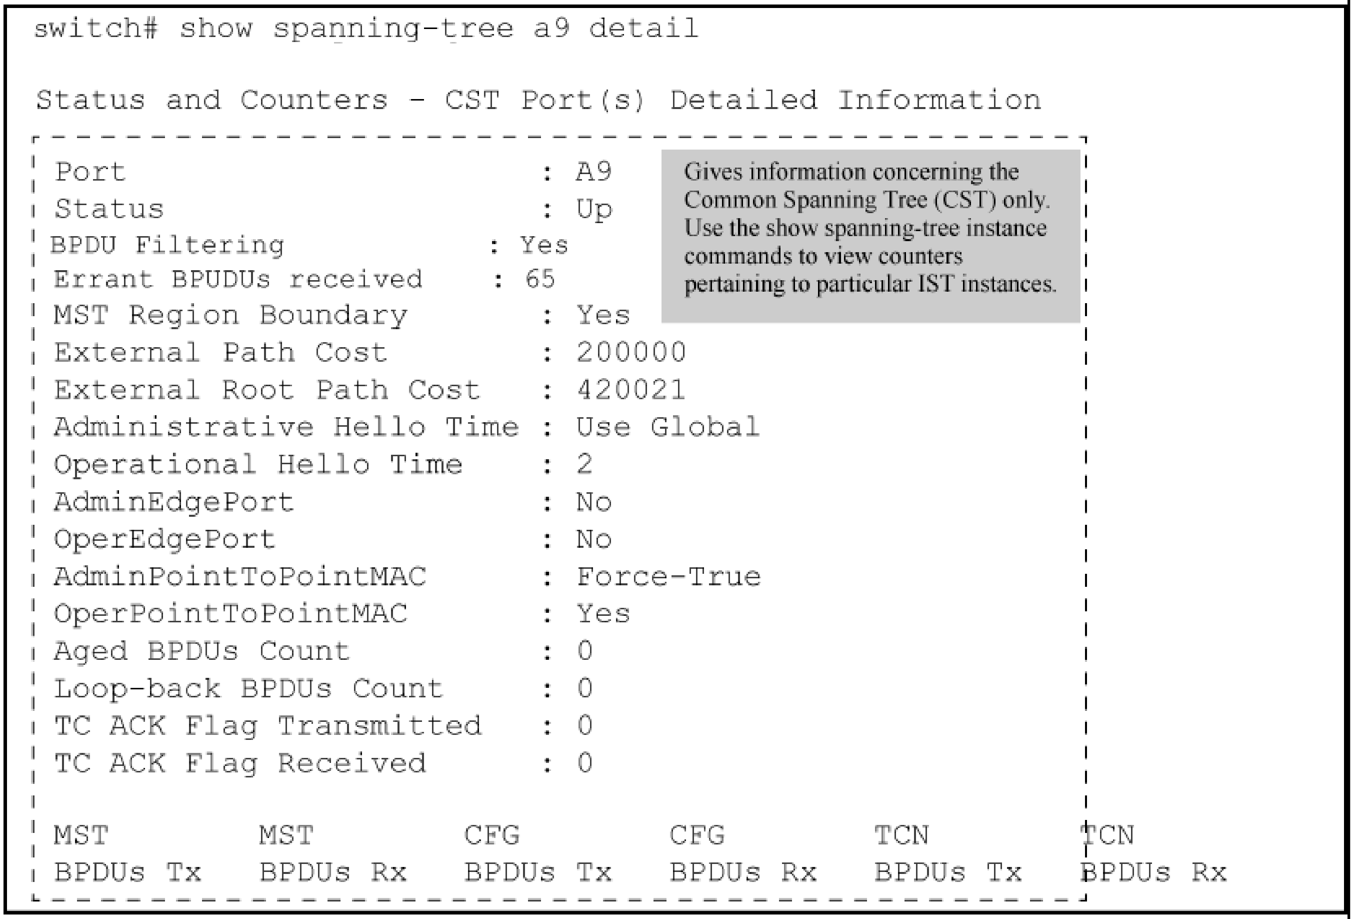

```
switch# show spanning-tree a9 detail

 Status and Counters - CST Port(s) Detailed Information

  Port                         : A9
  Status                       : Up
  BPDU Filtering               : Yes
  Errant BPUDUs received       : 65
  MST Region Boundary          : Yes
  External Path Cost           : 200000
  External Root Path Cost      : 420021
  Administrative Hello Time    : Use Global
  Operational Hello Time       : 2
  AdminEdgePort                : No
  OperEdgePort                 : No
  AdminPointToPointMAC         : Force-True
  OperPointToPointMAC          : Yes
  Aged BPDUs Count             : 0
  Loop-back BPDUs Count        : 0
  TC ACK Flag Transmitted      : 0
  TC ACK Flag Received         : 0

  MST         MST         CFG         CFG         TCN         TCN
  BPDUs Tx    BPDUs Rx    BPDUs Tx    BPDUs Rx    BPDUs Tx    BPDUs Rx
```

Gives information concerning the Common Spanning Tree (CST) only. Use the show spanning-tree instance commands to view counters pertaining to particular IST instances.

NOTE: This command gives information about the CST only. To view details of specific MST instances, use the `show spanning tree instance` commands.

Viewing status for a specific MST instance

The following commands display the MSTP statistics for a specified MST instance.

Syntax:

```
show spanning-tree instance [ist | 1..16]
```

Displays the MSTP statistics for either the IST instance or a numbered MST instance running on the switch.

Syntax:

```
show spanning-tree instance [ist | 1..16] detail
```

Displays status on all active ports for a specific instance of MSTP.

Syntax:

```
show spanning-tree port-list instance [ist | 1..16] detail
```

Displays status on specified ports for a specific instance of MSTP.

Viewing status for a specific instance of an MSTP

This shows how to display detailed status for all active ports for a specific instance of MSTP.

```
switch(config)# show spanning-tree instance 11
 MST Instance Information
  Instance ID : 11
  Mapped VLANs : 111,300
  Switch Priority         : 32768

  Topology Change Count    : 2
  Time Since Last Change   : 4 mins

 Regional Root MAC Address : 1cc1de-cfbc80
 Regional Root Priority    : 32768
 Regional Root Path Cost   : 400000
 Regional Root Port        : This switch is root
 Remaining Hops            : 20

                                                           Designated
 Port  Type      Cost      Priority Role       State      Bridge
 ----- --------- --------- -------- ---------- ---------- -------------
 1     10/100TX  200000    128      Root       Forwarding 1cc1de-cfbc80
 2     10/100TX  200000    128      Designated Forwarding 1cc1de-02a700
 3     10/100TX  Auto      112      Designated Forwarding 1cc1de-02a700
 4     10/100TX  Auto      128      Disabled   Disabled
 .        .       .         .          .         .
```

Viewing the MSTP configuration

MSTP configuration can be viewed at the global, per-instance, and regional level

Viewing the global MSTP configuration

This command displays the switch's basic and MST region spanning tree configuration, including basic port connectivity settings.

Syntax:

```
show spanning-tree config
```

The upper part of this output shows the switch's global spanning tree configuration that applies to the MST region. The port listing shows the spanning tree port parameter settings for the spanning tree region operation configured by the `spanning-tree port-list` command.

Syntax:

```
show spanning-tree port-list config
```

This command shows the same data as the above command, but lists the spanning tree port parameter settings for only the specified port or trunk. You can list data for a series of ports and port trunks by specifying the first and last port or trunk of any consecutive series of ports and trunks. For example, to display data for port A20-A24 andtrk1, use the command: `show spanning-tree a20-a24,trk1 config`

Figure 17: *Viewing the switch's global spanning tree configuration*

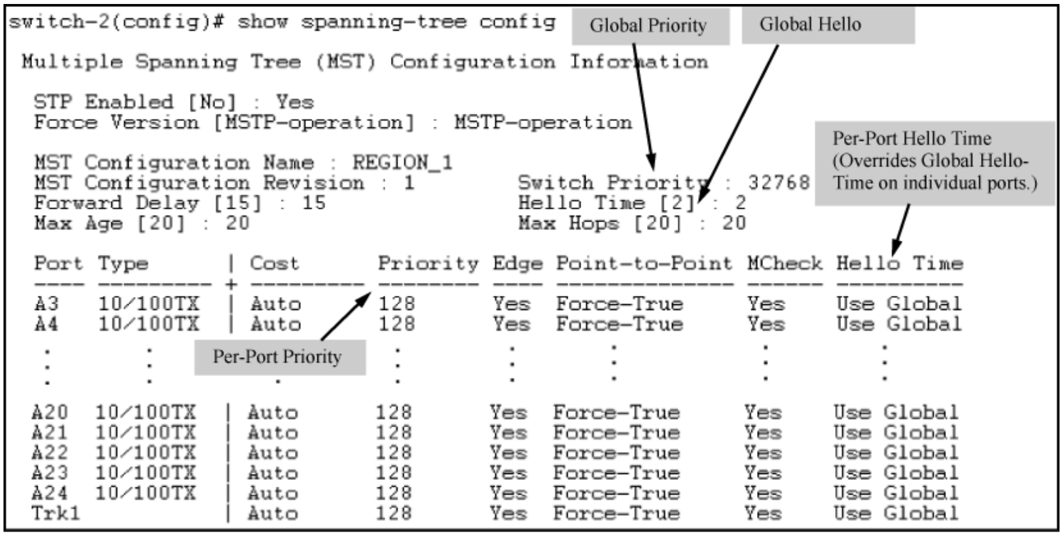

```
switch-2(config)# show spanning-tree config

 Multiple Spanning Tree (MST) Configuration Information

  STP Enabled [No] : Yes
  Force Version [MSTP-operation] : MSTP-operation

  MST Configuration Name : REGION_1
  MST Configuration Revision : 1       Switch Priority : 32768
  Forward Delay [15] : 15              Hello Time [2] : 2
  Max Age [20] : 20                    Max Hops [20] : 20

  Port Type      | Cost      Priority Edge Point-to-Point MCheck Hello Time
  ---- --------- + --------- -------- ---- -------------- ------ ----------
  A3   10/100TX  | Auto      128      Yes  Force-True     Yes    Use Global
  A4   10/100TX  | Auto      128      Yes  Force-True     Yes    Use Global
  .      .                   .        .      .            .        .
  .      .                   .        .      .            .        .
  .      .           .       .        .      .            .        .

  A20  10/100TX  | Auto      128      Yes  Force-True     Yes    Use Global
  A21  10/100TX  | Auto      128      Yes  Force-True     Yes    Use Global
  A22  10/100TX  | Auto      128      Yes  Force-True     Yes    Use Global
  A23  10/100TX  | Auto      128      Yes  Force-True     Yes    Use Global
  A24  10/100TX  | Auto      128      Yes  Force-True     Yes    Use Global
  Trk1           | Auto      128      Yes  Force-True     Yes    Use Global
```

Global Priority
Global Hello
Per-Port Hello Time (Overrides Global Hello-Time on individual ports.)
Per-Port Priority

Viewing per-instance MSTP configurations

These commands display the per-instance port configuration and current state, along with instance identifiers and regional root data.

Syntax:

```
show spanning-tree config instance [ist | 1..16]
```

The upper part of this output shows the instance data for the ist or for the specified instance. The lower part of the output lists the spanning tree port settings for the specified instance.

Syntax:

```
show spanning-tree port-list config instance [ist | 1..16]
```

This command shows the same data as the preceding command, but lists the spanning tree port parameter settings for only the specified port or trunk. You can list data for a series of ports and port trunks by specifying the first and last port or trunk of any consecutive series of ports and trunks.

Viewing port data

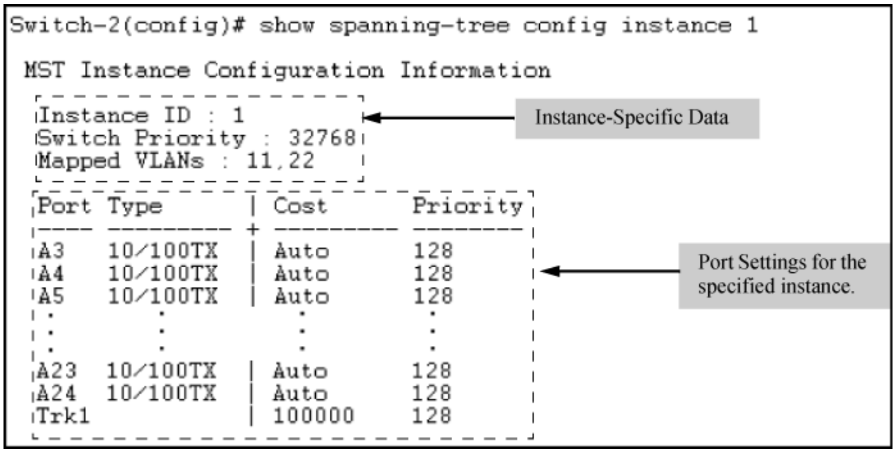

```
Switch-2(config)# show spanning-tree config instance 1

 MST Instance Configuration Information

  Instance ID : 1
  Switch Priority : 32768
  Mapped VLANs : 11,22

  Port Type       | Cost      Priority
  ---- --------- + --------- --------
  A3   10/100TX  | Auto      128
  A4   10/100TX  | Auto      128
  A5   10/100TX  | Auto      128
  .     .          .          .
  .     .          .          .
  .     .          .          .
  A23  10/100TX  | Auto      128
  A24  10/100TX  | Auto      128
  Trk1           | 100000    128
```

To display data for ports A20-A24 and trk1, you would use the command:

```
switch(config)# show spanning-tree a20-a24,trk1 config instance 1
```

Viewing the region-level configuration

This command is useful for quickly verifying the allocation of VLANs in the switch's MSTP configuration, and for viewing the configured region identifiers.

Syntax:

```
show spanning-tree mst-config
```

NOTE:

The switch computes the MSTP Configuration Digest from the VID to MSTI configuration mappings on the switch itself. As required by the 802.1s standard, all MSTP switches within the same region must have the same VID to MSTI assignments, and any given VID can be assigned to either the IST or one of the MSTIs within the region. Thus, the MSTP Configuration Digest must be identical for all MSTP switches intended to belong to the same region. When comparing two MSTP switches, if their Digest identifiers do not match, they cannot be members of the same region.

Viewing a region-level configuration

```
switch(config)# show spanning-tree net-config

 MST Configuration Identifier Information

  MST Configuration Name : REGION_1
  MST Configuration Revision : 1
  MST Configuration Digest : 0xDAD6A13EC5141980B7EBDA71D8991E7C

  IST Mapped VLANs : 1,66

  Instance ID Mapped VLANs
  -------- ---------------
  1        11,22
  2        33,44,55
```

Viewing the pending MSTP configuration

This command displays the MSTP configuration the switch will implement if you execute the `spanning tree pending apply` command.

Syntax:

```
show spanning-tree pending [instance | mst-config]
```

instance [*1..16* | ist]

Lists region, instance ID, and VLAN information for the specified, pending instance.

mst-config

Lists region, IST instance VLANs, numbered instances, and assigned VLAN information for the pending MSTP configuration.

Viewing a pending configuration

```
switch(config)# show spanning-tree pending instance 3

 Pending MST Instance Configuration Information

  MST Configuration Name : New-Version_01
  MST Configuration Revision : 1
  Instance ID : 3
  Mapped VLANs : 3

switch(config)# show spanning-tree pending mst-config

 Pending MST Configuration Identifier Information

  MST Configuration Name : New-Version_01
  MST Configuration Revision : 1

  IST Mapped VLANs : 1,2,4-4094

  Instance ID Mapped VLANs
  ----------- ---------------------------------------
  3           3
```

MSTP operating rules

- All switches in a region must be configured with the same set of VLANs, the same MST configuration name and MST configuration number.
- Within a region, a VLAN can be allocated to either a single MSTI or to the region's IST instance.
- All switches in a region must have the same VID-to-MST instance assignment.
- There is one root MST switch per configured MST instance.
- Because boundary ports provide the VLAN connectivity between regions, all boundary ports on a region's root switch should be configured as members of all static VLANs defined in the region.
- There is one root switch for the Common and Internal Spanning Tree (CIST). At any given time, all switches in the network will use the per-port `hello-time` parameter assignments configured on the CIST root switch.

- Where multiple MST regions exist in a network, there is only one active, physical communication path between any two regions, or between an MST region and an STP or RSTP switch. MSTP blocks any other physical paths as long as the currently active path remains in service.
- Within a network, an MST region appears as a virtual RSTP bridge to other spanning tree entities (other MST regions, and any switches running 802.1D or 802.1w spanning tree protocols).
- Within an MSTI, there is one physical communication path between any two nodes, regardless of how many VLANs belong to the MSTI. Within an IST instance, there is also one spanning tree across all VLANs belonging to the IST instance.
- An MSTI comprises a unique set of VLANs and forms a single spanning tree instance within the region to which it belongs.
- A dynamic VLAN learned by GVRP will always be placed in the IST instance and cannot be moved to any configured MST instance.Starting in software release 13.X.X, dynamically learned GVRP VLANs can be mapped to MSTIs and support MSTP load balancing.
- In software release 13.X.X and later, you can preconfigure static and dynamic VLAN ID-to-MSTI mappings before the VLAN is created on the switch. Later, when the static VLAN ID is configured or a dynamic GVRP VLAN is learned, the VLAN is automatically associated with the preconfigured MSTI.
- Communication between MST regions uses a single spanning tree.
- If a port on a switch configured for MSTP receives a legacy (STP/802.1D or RSTP/802.1w) BPDU, it automatically operates as a legacy port. In this case, the MSTP switch interoperates with the connected STP or RSTP switch as a separate MST region.
- Within an MST region, there is one logical forwarding topology per instance, and each instance comprises a unique set of VLANs. Where multiple paths exist between a pair of nodes using VLANs belonging to the same instance, all but one of those paths will be blocked for that instance. However, if there are different paths in different instances, all such paths are available for traffic. Separate forwarding paths exist through separate spanning tree instances.
- A port can have different states (forwarding or blocking) for different instances (which represent different forwarding paths).
- MSTP interprets a switch mesh as a single link.

Troubleshooting an MSTP configuration

Cause

This section describes the `show spanning-tree` commands to use to monitor, troubleshoot, and debug the operation of a multiple-instance spanning tree configuration in a network.

The `show spanning-tree` commands described in this section allow for focusing on increasingly specific levels of operation. For example, you can display debug information for:

- All MST instances
- All ports used in one MST instance
- A specific port or several ports used in one MST instance

Also, you can display the change history for the root (bridge) switch used as the single forwarding path for:

- All MST regions, STP bridges, and RSTP bridges in an STP network
- All VLANs on MSTP switches in a region
- All VLANs on MSTP switches in an mst instance

Viewing the change history of root bridges

The `show spanning-tree root-history` command allows you to display change history information (up to 10 history entries) for a specified root bridge in any of the following MSTP topologies:

- Common Spanning Tree (`cst`):Provides connectivity in a bridged network between MST regions, STP LANs, and RSTP LANs.
- Internal Spanning Tree (`ist`):Provides connectivity within an MST region for VLANs associated with the default Common and Internal Spanning Tree (CIST) instance in your network (VLANs that have not been mapped to an MST instance).
- MST Instance (`mst`):Connects all static and (from release 13.X.X) dynamic VLANs assigned to a multiple spanning tree instance.

Syntax:

```
show spanning tree root-history [cst | ist | mst] instance-id
```

Displays the change history for the root bridge in the specified MSTP topology.

`cst`

Displays the change history for the root bridge of a spanning tree network, including MST regions and STP and RSTP bridges.

`ist`

Displays the change history for the root bridge in the IST instance of an MST region.

`mst` ***instance-id***

Displays the change history for the root bridge in an MST instance, where *instance-id* is an ID number from 1 to 16.

Use the `show spanning-tree root-history` command to view the number and dates of changes in the assignment of a root bridge. Possible intrusion into your MST network may occur if an unauthorized external device gains access to a spanning tree by posing as the root device in a topology. To prevent an MST port connected to the device from being selected as the root port in a topology, use the `spanning-tree root-guard` command.

Sample output of the `show spanning-tree root-history` command for different MSTP topologies

The following examples show sample output of the `show spanning-tree root-history` command for different MSTP topologies. In each example, the root bridge ID is displayed in the format: *priority*: *mac-address*

Where:

- *priority*

  is the MSTP switch priority calculated for one of the following:

    - The IST (regional) root switch using the `spanning-tree priority` command
    - An MSTI root switch using the `spanning-tree instance priority` command
- *mac-address*

  is the MAC address of the root (bridge) switch.

Viewing `show spanning-tree root-history` CST output

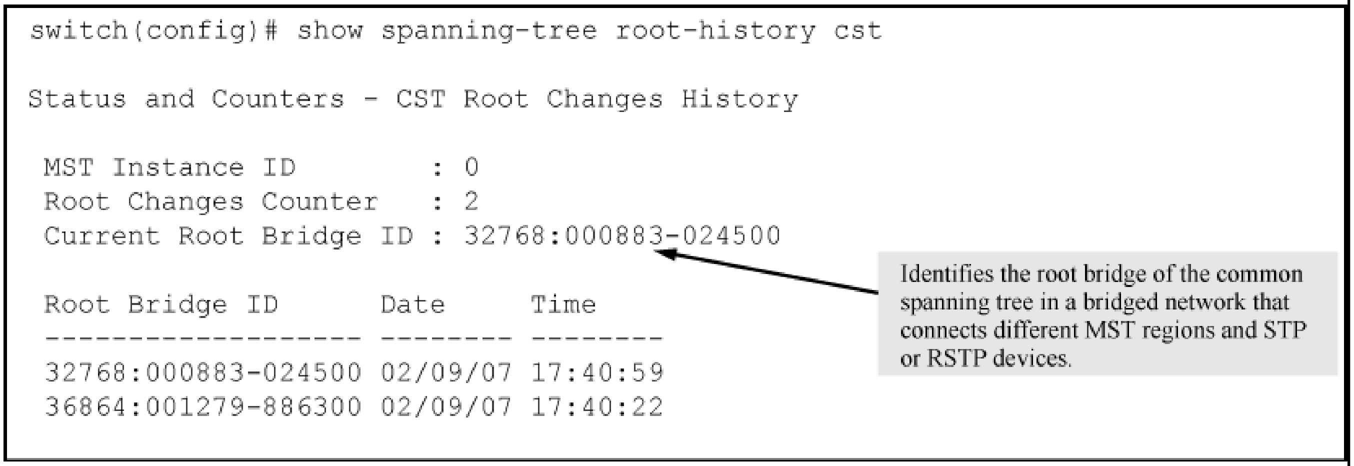

```
switch(config)# show spanning-tree root-history cst

Status and Counters - CST Root Changes History

 MST Instance ID         : 0
 Root Changes Counter    : 2
 Current Root Bridge ID : 32768:000883-024500

 Root Bridge ID       Date     Time
 ------------------- -------- --------
 32768:000883-024500 02/09/07 17:40:59
 36864:001279-886300 02/09/07 17:40:22
```

Identifies the root bridge of the common spanning tree in a bridged network that connects different MST regions and STP or RSTP devices.

Viewing `show spanning-tree root-history` IST output

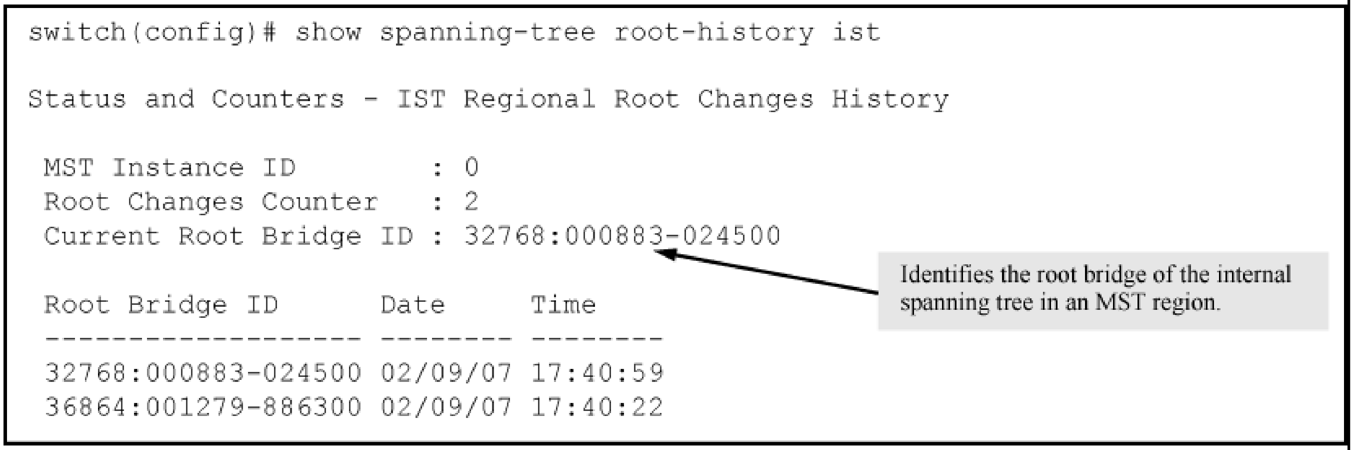

```
switch(config)# show spanning-tree root-history ist

Status and Counters - IST Regional Root Changes History

 MST Instance ID         : 0
 Root Changes Counter    : 2
 Current Root Bridge ID : 32768:000883-024500

 Root Bridge ID       Date     Time
 ------------------- -------- --------
 32768:000883-024500 02/09/07 17:40:59
 36864:001279-886300 02/09/07 17:40:22
```

Identifies the root bridge of the internal spanning tree in an MST region.

Viewing `show spanning-tree root-history` MSTI output

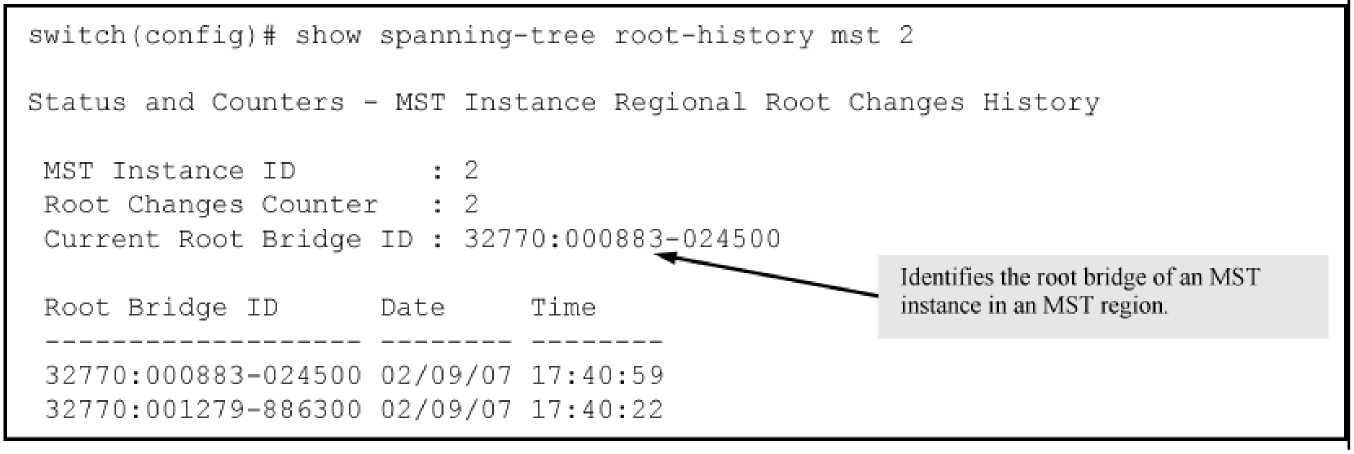

```
switch(config)# show spanning-tree root-history mst 2

Status and Counters - MST Instance Regional Root Changes History

 MST Instance ID         : 2
 Root Changes Counter    : 2
 Current Root Bridge ID : 32770:000883-024500

 Root Bridge ID       Date     Time
 -------------------- -------- --------
 32770:000883-024500 02/09/07 17:40:59
 32770:001279-886300 02/09/07 17:40:22
```

Identifies the root bridge of an MST instance in an MST region.

Enabling traps and viewing trap configuration

Syntax

`spanning-tree trap {errant-bpdu | loop-guard | new-root | root-guard}`

`no spanning-tree trap {errant-bpdu | loop-guard | new-root | root-guard}`

Enables SNMP traps. The `no` form of the command disables SNMP traps.

Syntax

`show spanning-tree traps`

Displays the current spanning tree trap configuration on the switch.

Viewing spanning tree traps in their default configuration

```
switch# show spanning-tree traps

 Status and Counters - STP Traps Information

 Trap Name                | Status
 ----------------------- + --------
 errant-bpdu              | Disabled
 new-root                 | Disabled
 root-guard               | Disabled
 loop-guard               | Disabled
```

Viewing debug counters for all MST instances

The `show spanning-tree debug-counters` command allows you to display the aggregate values of all MSTP debug counters that are maintained on a switch. These aggregate values are a summary of the information collected from all ports and from all spanning tree instances that forward traffic on switch ports.

Use the displayed diagnostic information to globally monitor MSTP operation on a per-switch basis.

Syntax:

`show spanning-tree debug-counters`

Displays debug counters for MSTP activity on all ports configured for VLANs used in spanning tree instances.

Viewing output for debug counters

The following example shows sample output of the `show spanning-tree debug-counters` command for all ports.

```
switch(config)# show spanning-tree debug-counters

 Status and Counters - MSTP Bridge Common Debug Counters Information

  Counter Name                       Aggregated Value Collected From
  ---------------------------------  ---------------- --------------
  Invalid BPDUs                      0                CIST
  Errant BPDUs                       170927           CIST
  MST Config Error BPDUs             0                CIST
  Looped-back BPDUs                  0                CIST
  Starved BPDUs/MSTI MSGs            0                CIST/MSTIs
  Exceeded Max Age BPDUs             0                CIST
  Exceeded Max Hops BPDUs/MSTI MSGs  0                CIST/MSTIs
  Topology Changes Detected          2                CIST/MSTIs
  Topology Changes Tx                6                CIST/MSTIs
  Topology Changes Rx                4                CIST/MSTIs
  Topology Change ACKs Tx            0                CIST
  Topology Change ACKs Rx            0                CIST
  TCN BPDUs Tx                       0                CIST
  TCN BPDUs Rx                       0                CIST
  CFG BPDUs Tx                       0                CIST
  CFG BPDUs Rx                       0                CIST
  RST BPDUs Tx                       0                CIST
  RST BPDUs Rx                       0                CIST
  MST BPDUs/MSTI MSGs Tx             10               CIST/MSTIs
  MST BPDUs/MSTI MSGs Rx             341802           CIST/MSTIs
```

Viewing debug counters for one MST instance

The `show spanning-tree debug-counters instance` command allows you to display the aggregate values of all MSTP debug counters maintained on a switch for a specified spanning tree instance. These aggregate values are a summary of information collected from all ports that have VLANs assigned to the specified instance.

Use this command to troubleshoot the global MSTP diagnostic information displayed in `show spanning-tree debug-counters` command output when you suspect unauthorized MSTP activity in a specific MST instance.

Syntax:

```
show spanning-tree debug-counters instance instance-id
```

Displays debug counters for MSTP activity on all ports configured for VLANs in the specified MST instance.

The valid values for `instance` *instance-id* are 0 to 16, where 0 specifies the default MST (CIST) instance and 1 to 16 specify a multiple spanning tree (MST) instance.

Viewing bug counters for a CIST instance

The following example shows sample output of the `show spanning-tree debug-counters instance` command when applied to the Common and Internal Spanning Tree (CIST) instance (default MST instance 0) in the network.

```
switch(config)# show spanning-tree debug-counters instance 0
```

```
Status and Counters - CIST Common Debug Counters Information

 MST Instance ID : 0

 Counter Name                      Aggregated Value Collected From
 --------------------------------- ---------------- ---------------
 Invalid BPDUs                     0                Ports
 Errant BPDUs                      172603           Ports
 MST Config Error BPDUs            0                Ports
 Looped-back BPDUs                 0                Ports
 Starved BPDUs                     0                Ports
 Exceeded Max Age BPDUs            0                Ports
 Exceeded Max Hops BPDUs           0                Ports
 Topology Changes Detected         1                Ports
 Topology Changes Tx               3                Ports
 Topology Changes Rx               2                Ports
 Topology Change ACKs Tx           0                Ports
 Topology Change ACKs Rx           0                Ports
 TCN BPDUs Tx                      0                Ports
 TCN BPDUs Rx                      0                Ports
 CFG BPDUs Tx                      0                Ports
 CFG BPDUs Rx                      0                Ports
 RST BPDUs Tx                      0                Ports
 RST BPDUs Rx                      0                Ports
 MST BPDUs Tx                      5                Ports
 MST BPDUs Rx                      172577           Ports
```

Viewing debug counters for ports in an MST instance

The `show spanning-tree debug-counters instance ports` command displays the aggregate values of all MSTP debug counters maintained on one or more ports used by a specified spanning tree instance. These aggregate values are a summary of information collected from the specified ports that have VLANs assigned to the specified instance.

Use this command to troubleshoot at a finer level the more general MSTP diagnostic information displayed in the `show spanning-tree debug-counters instance` command output, when you suspect unauthorized MSTP activity on one or more MST ports in an MST instance.

Syntax:

```
show spanning-tree debug-counters instance instance-id ports port-list
```

Displays debug counters for MSTP activity on the specified ports configured for VLANs in the specified MST instance.

`instance` *instance-id*

The valid values for *instance-id* are from 0 to 16, where 0 specifies the default MST (CIST) instance and 1 to 16 specify an MST instance.

`ports` *port-list*

Specifies one or more MST ports or trunk ports. In the port list, enter a series of ports by separating the first and last ports in the series with a dash (-); for example, `a2-a8` or `trk1-trk3`. Separate individual ports and series of ports with a comma; for example, `a2-a8, a20, trk1, trk4-trk5`.

Viewing debug counters for a CIST and MST instance

The following example shows sample output of the `show spanning-tree debug-counters instance ports` command for both the CIST (default MST instance 0) and an MST instance (instance 2) on port A15.

```
switch(config)# show spanning-tree debug-counters instance 0 ports a15

 Status and Counters - CIST Port(s) Debug Counters Information

  MST Instance ID : 0
  Port : A15

  Counter Name                Value      Last Updated
  --------------------------- ---------- -----------------
  Invalid BPDUs               0
  Errant BPDUs                0
  MST Config Error BPDUs      0
  Looped-back BPDUs           0
  Starved BPDUs               0
  Exceeded Max Age BPDUs      0
  Exceeded Max Hops BPDUs     0
  Topology Changes Detected   1         02/09/07 17:40:59
  Topology Changes Tx         3         02/09/07 17:41:03
  Topology Changes Rx         2         02/09/07 17:41:01
  Topology Change ACKs Tx     0
  Topology Change ACKs Rx     0
  TCN BPDUs Tx                0
  TCN BPDUs Rx                0
  CFG BPDUs Tx                0
  CFG BPDUs Rx                0
  RST BPDUs Tx                0
  RST BPDUs Rx                0
  MST BPDUs Tx                5         02/09/07 17:41:03
  MST BPDUs Rx                173540    02/13/07 18:05:34
```

Viewing debug counters output for one port in an MST instance

The following example shows spanning tree debug-counters instance ports command output for one port in an MST instance.

```
switch(config)# show spanning-tree debug-counters instance 2 ports a15

 Status and Counters - MSTI Port(s) Debug Counters Information

   MST Instance ID : 2
   Port : A15

   Counter Name                Value      Last Updated
   --------------------------- ---------- -----------------
   Starved MSTI MSGs           0
   Exceeded Max Hops MSTI MSGs 0
   Topology Changes Detected   1          02/09/07 17:40:59
   Topology Changes Tx         3          02/09/07 17:41:03
   Topology Changes Rx         2          02/09/07 17:41:01
   MSTI MSGs Tx                5          02/09/07 17:41:03
   MSTI MSGs Rx                173489     02/13/07 18:03:52
```

Field descriptions in MSTP debug command output

The following table contains descriptions of the debugging information displayed in the output of `show spanning-tree debug-counters` commands.

Table 14: *MSTP debug command output: field descriptions*

Field	Displays the number of...
`Invalid BPDUs`	Received BPDUs that failed standard MSTP (802.1Q-REV/D5.0 14.4) validation checks and were dropped. This counter is maintained by the CIST (default MST instance 0) on a per-port basis.
`Errant BPDUs`	Received BPDUs that were dropped on a port that is configured to not expect BPDU packets. This counter is maintained by the CIST (MST instance, 0default MST instance 0 in the network) on a per-port basis and is incremented each time a BPDU packet is received on a port configured with the BPDU filter to ignore incoming BPDU packets (`spanning-tree bpdu-filter` command) or the BPDU protection feature to disable the port when BPDU packets are received (`spanning-tree bpdu-protection` command).
`MST Config Error BPDUs`	BPDUs received from a neighbor bridge with inconsistent MST configuration information. For example, BPDUs from a transmitting bridge may contain the same MST configuration identifiers (region name and revision number) and format selector as the receiving bridge, but the value of the Configuration Digest field (VLAN ID assignments to regional IST and MST instances) is different. This difference indicates a probable configuration error in MST region settings on the communicating bridges. The received BPDU is still processed by MSTP. This counter is maintained by the CIST (default MST instance 0) on a per-port basis.
`Looped-back BPDUs`	Times a port has received self-sent BPDU packets as the result of an external loop condition in which the BPDUs were looped back to the originating transmission port. The received BPDU is still processed by MSTP and the port changes to a blocked state. This counter is maintained by the CIST (default MST instance 0) on a per-port basis.
`Starved BPDUs`	Times that no BPDUs are received within the scheduled interval (three times the Hello Time value configured with the `spanning-tree hello-time` command) from a downstream CIST-designated peer port on the CIST root, alternate, or backup port. As a result, the "starved" port triggers a spanning tree topology regeneration. This counter is maintained by the CIST (default MST instance 0) on a per-port basis.

Table Continued

Field	Displays the number of...
`Starved MSTI MSGs`	Times that no BPDUs are received within the scheduled interval (three times the Hello Time value configured with the `spanning-tree hello-time` command) from a downstream MSTI-designated peer port on the MSTI root, alternate, or backup port. As a result, the "starved" port triggers a spanning tree topology regeneration. This counter is maintained by the CIST (default MST instance 0) on a per-port basis.
`Exceeded Max Age BPDUs`	Times that a BPDU packet is received from a bridge external to the MST region with a Message Age value greater than the configured value of the Max Age parameter (`spanning-tree maximum age` command). This may occur if the receiving bridge is located too far from the root bridge (beyond the configured size of the spanning tree domain on the root bridge) or if a BPDU packet with invalid root information is continuously circulating between bridges in a spanning tree domain and needs to be aged out. This counter is maintained by the CIST (default MST instance 0) on a per-port basis.
`Exceeded Max Hops BPDUs`	Times that a BPDU packet is received from a bridge internal to the MST region with a CIST Remaining Hops value less than or equal to 1. This may occur if the receiving bridge is located too far from the CIST regional root bridge (beyond the configured size of the MST region on the CIST regional root bridge) or if a BPDU packet with invalid CIST regional root bridge information is continuously circulating between bridges in the MST Region and needs to be aged out. This counter is maintained by the CIST (default MST instance 0 in the region) on a per-port basis.
`Exceeded Max Hops MSTI MSGs`	Times that an MSTI MSG packet is received from a bridge internal to the MST region with an MSTI Remaining Hops value less than or equal to 1. This may occur if the receiving bridge is located too far from the MSTI regional root bridge (beyond the configured size of the MST region on the MSTI regional root bridge) or if a BPDU packet with invalid MSTI regional root bridge information is continuously circulating between bridges in an MST region and needs to be aged out. This counter is maintained on a per-MSTI per-port basis.
`Topology Changes Detected`	Times that a Topology Change event is detected by the CIST or MSTI port and the port triggers a topology change propagation throughout the network. A Topology Change event occurs when a non-edge port enters forwarding state. This counter is maintained on a per-CIST per-port and on a per-MSTI per-port basis.

Table Continued

Field	Displays the number of...
`Topology Changes Tx`	Times that Topology Change information is propagated (sent out) through the port to the rest of the network.For a CIST port, the counter is the number of times that a CFG, RST, or MST BPDU with the TC flag set is transmitted out of the port.For an MSTI port, the counter is the number of times that an MSTI configuration message with the TC flag set is transmitted out of the port. This counter is maintained on a per-CIST per-port and on a per-MSTI per-port bases.
`Topology Changes Rx`	Times that Topology Change information is received from the peer port.For a CIST port, the counter is the number of times that a CFG, RST, or MST BPDU with the TC flag set is received.For an MSTI port, the counter is the number of times that an MSTI configuration message with the TC flag set is received. This counter is maintained on a per-CIST per-port and on a per-MSTI per-port basis.
`Topology Change ACKs Tx`	Times that the Topology Change acknowledgement is transmitted through the port (number of CFG, RST or MST BPDUs transmitted with the Topology Change Acknowledge flag set). This counter is maintained by the CIST (default MST instance 0) on a per-port basis.
`Topology Change ACKs Rx`	Times the Topology Change acknowledgement is received on the port (number of CFG, RST or MST BPDUs received with the Topology Change Acknowledge flag set). This counter is maintained by the CIST (default MST instance 0) on a per-port basis.
`TCN BPDUs Tx`	Topology Change Notification BPDUs that are transmitted through the port. This counter is maintained by the CIST (default MST instance 0) on a per-port basis.
`TCN BPDUs Rx`	Topology Change Notification BPDUs that are received on the port. This counter is maintained by the CIST (default MST instance 0) on a per-port basis.
`CFG BPDUs Tx`	802.1D Configuration BPDUs that are transmitted through the port. This counter is maintained by the CIST (default MST instance 0) on a per-port basis.
`CFG BPDUs Rx`	802.1D Configuration BPDUs that are received on the port. This counter maintained by the CIST (default MST instance 0) on a per-port basis.

Table Continued

Field	Displays the number of...
`RST BPDUs Tx`	802.1w RST BPDUs that are transmitted through the port. This counter is maintained by the CIST (default MST instance 0) on a per-port basis.
`RST BPDUs Rx`	802.1w RST BPDUs that are received on the port. This counter is maintained by the CIST (default MST instance 0) on a per-port basis.
`MST BPDUs Tx`	802.1s MST BPDUs that are transmitted through the port. This counter is maintained by the CIST (default MST instance 0) on a per-port basis.
`MST BPDUs Rx`	802.1s MST BPDUs that are received on the port. This counter is maintained by the CIST (default MST instance 0) on a per-port basis.
`MSTI MSGs Tx`	Times that a configuration message for a specific MSTI was encoded in (802.1s) MST BPDUs that are transmitted through the port. This counter is maintained on a per-MSTI per-port basis.
`MSTI MSGs Rx`	Times that the MSTI detected a configuration message destined to the MSTI in (802.1s) MST BPDUs received on the port. This counter is maintained on a per-MSTI per-port basis.

Troubleshooting MSTP operation

Table 15: *Troubleshooting MSTP operation*

Problem	Possible cause
Duplicate packets on a VLAN, or packets not arriving on a LAN at all.	The allocation of VLANs to MSTIs may not be identical among all switches in a region.
A switch intended to operate in a region does not receive traffic from other switches in the region.	An MSTP switch intended for a particular region may not have the same configuration name or region revision number as the other switches intended for the same region. The MSTP configuration name (`spanning-tree config-name` command) and MSTP configuration revision number (`spanning-tree config-revision` command) must be identical on all MSTP switches intended for the same region. Another possible cause is that the set of VLANs and VLAN ID-to-MSTI mappings (`spanning-tree instance vlan` command) configured on the switch may not match the set of VLANs and VLAN ID-to-MSTI mappings configured on other switches in the intended region.

BPDU

BPDUs are data messages that are exchanged across the switches within an extended LAN that uses a spanning tree protocol topology. BPDU packets contain information on ports, addresses, priorities, and costs.

About BPDU protection

BPDU protection is a security feature designed to protect the active STP topology by preventing spoofed BPDU packets from entering the STP domain. In a typical implementation, BPDU protection would be applied to edge ports connected to end user devices that do not run STP. If STP BPDU packets are received on a protected port, the feature will disable that port and alert the network manager via an SNMP trap as shown below.

BPDU protection enabled at the network edge

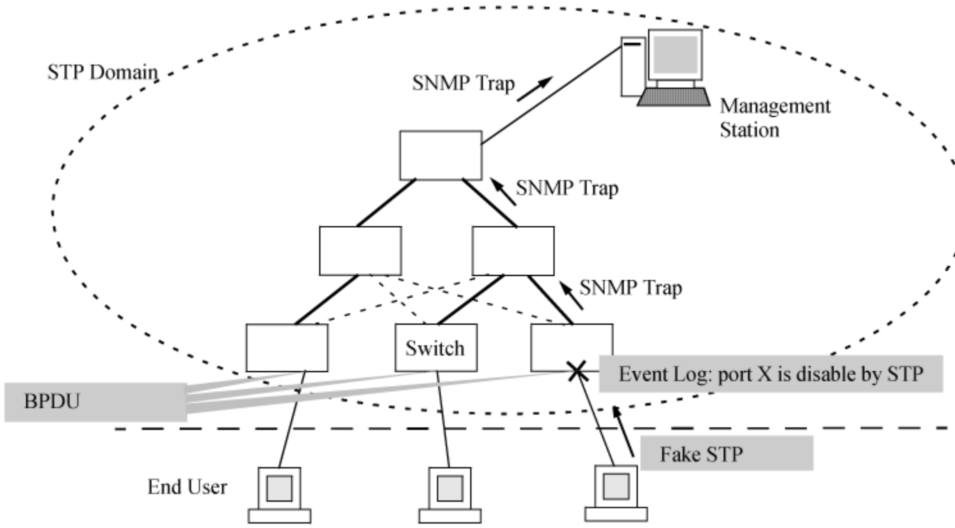

Viewing BPDU protection status

Syntax:

```
show spanning-tree bpdu-protection
```

Displays a summary listing of ports with BPDU protection enabled. To display detailed per port status information, enter the specific port numbers as shown here.

Figure 18: *Viewing BPDU protection status*

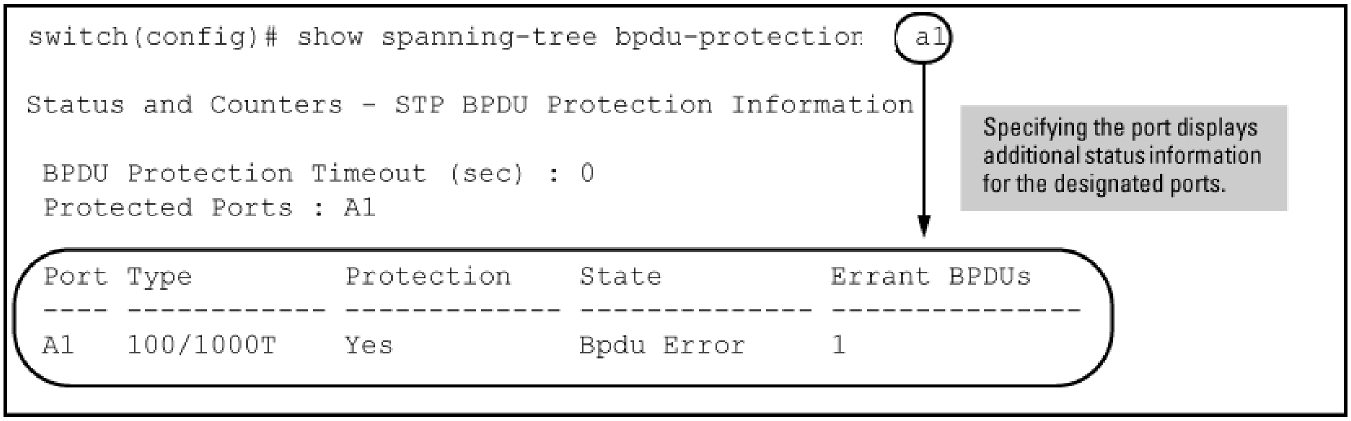

```
switch(config)# show spanning-tree bpdu-protection a1

Status and Counters - STP BPDU Protection Information

 BPDU Protection Timeout (sec) : 0
 Protected Ports : A1

 Port Type          Protection     State          Errant BPDUs
 ---- ------------ -------------- -------------- ----------------
 A1   100/1000T     Yes            Bpdu Error     1
```

Specifying the port displays additional status information for the designated ports.

BPDU protected ports are displayed as separate entries of the spanning tree category within the configuration file.

Figure 19: *Viewing BPDU filters using the* `show configuration` *command*

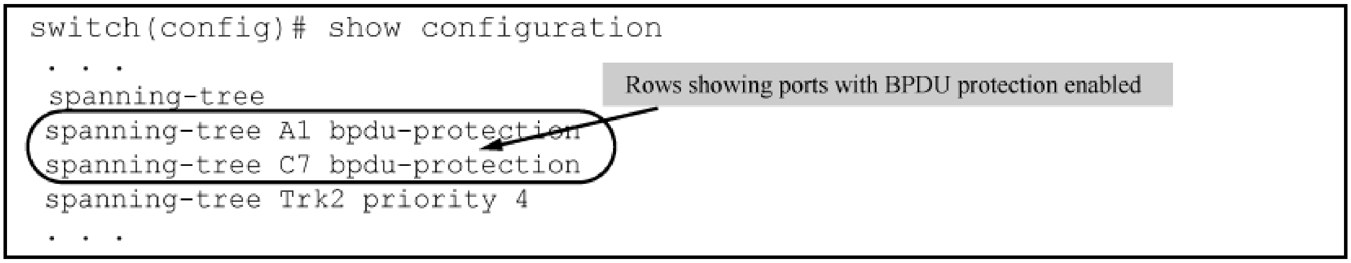

```
switch(config)# show configuration
 . . .
 spanning-tree
spanning-tree A1 bpdu-protection
spanning-tree C7 bpdu-protection
spanning-tree Trk2 priority 4
 . . .
```

Configuring BPDU filtering

The STP BPDU filter feature allows control of spanning tree participation on a per-port basis. It can be used to exclude specific ports from becoming part of spanning tree operations. A port with the BPDU filter enabled will ignore incoming BPDU packets and stay locked in the spanning tree forwarding state. All other ports will maintain their role.

Syntax:

```
spanning-tree [port-list | all] bpdu-filter
```

```
no spanning-tree [port-list | all] bpdu-filter
```

Enables or disables the BPDU filter feature on specified port(s). This forces a port to always stay in the forwarding state and be excluded from standard STP operation.

Sample scenarios in which this feature may be used are:

- To have STP operations running on selected ports of the switch rather than every port of the switch at a time.
- To prevent the spread of errant BPDU frames.
- To eliminate the need for a topology change when a port's link status changes. For example, ports that connect to servers and workstations can be configured to remain outside of spanning tree operations.
- To protect the network from denial of service attacks that use spoofing BPDUs by dropping incoming BPDU frames. For this scenario, BPDU protection offers a more secure alternative, implementing port shut-down and a detection alert when errant BPDU frames are received.

CAUTION: Ports configured with the BPDU filter mode remain active (learning and forward frames); however, spanning tree cannot receive or transmit BPDUs on the port. The port remains in a forwarding state, permitting all broadcast traffic. This can create a network storm if there are any loops (that is, trunks or redundant links) using these ports. If you suddenly have a high load, disconnect the link and disable the bpdu-filter (using the `no` command).

Configuring BPDU filtering

To configure BPDU filtering on port a9, enter:

```
switch(config)# spanning-tree a9 bpdu-filter
```

Viewing BPDU filtering

Syntax:

`spanning-tree show` *`port`* `configuration`

Displays the BPDU filter state.

Viewing BPDU filter status using the `show spanning tree` command

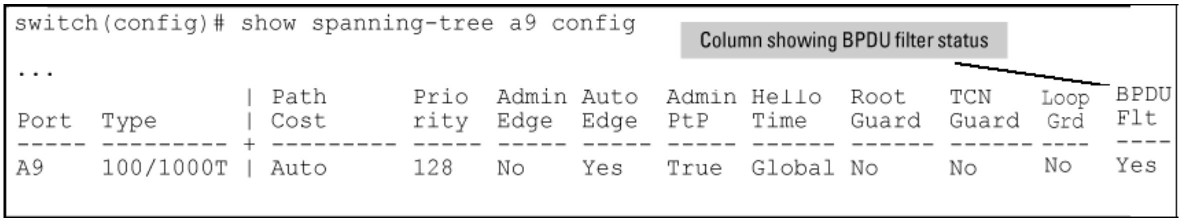

```
switch(config)# show spanning-tree a9 config

...
                 | Path       Prio  Admin Auto  Admin Hello  Root   TCN    Loop BPDU
Port  Type       | Cost       rity  Edge  Edge  PtP   Time   Guard  Guard  Grd  Flt
----- ---------- + ---------- ----- ----- ----- ----- ------ ------ ------ ---- ----
A9    100/1000T  | Auto       128   No    Yes   True  Global No     No     No   Yes
```

Viewing BPDU filters using the `show configuration` command

BPDU filters per port are displayed as separate entries of the spanning tree category within the configuration file.

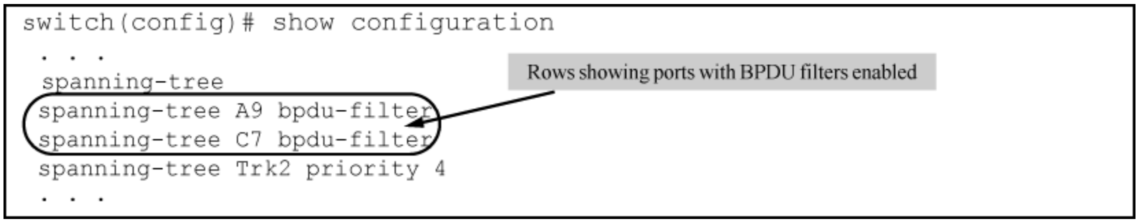

```
switch(config)# show configuration
  . . .
  spanning-tree
  spanning-tree A9 bpdu-filter
  spanning-tree C7 bpdu-filter
  spanning-tree Trk2 priority 4
  . . .
```

Configuring and managing BPDU protection

BPDU protection is a security feature designed to protect the active STP topology by preventing spoofed BPDU packets from entering the STP domain. In a typical implementation, BPDU protection would be applied to edge ports connected to end user devices that do not run STP. If STP BPDU packets are received on a protected port, the feature will disable that port and alert the network manager via an SNMP trap as shown in the following diagram.

Figure 20: *BPDU protection enabled at the network edge*

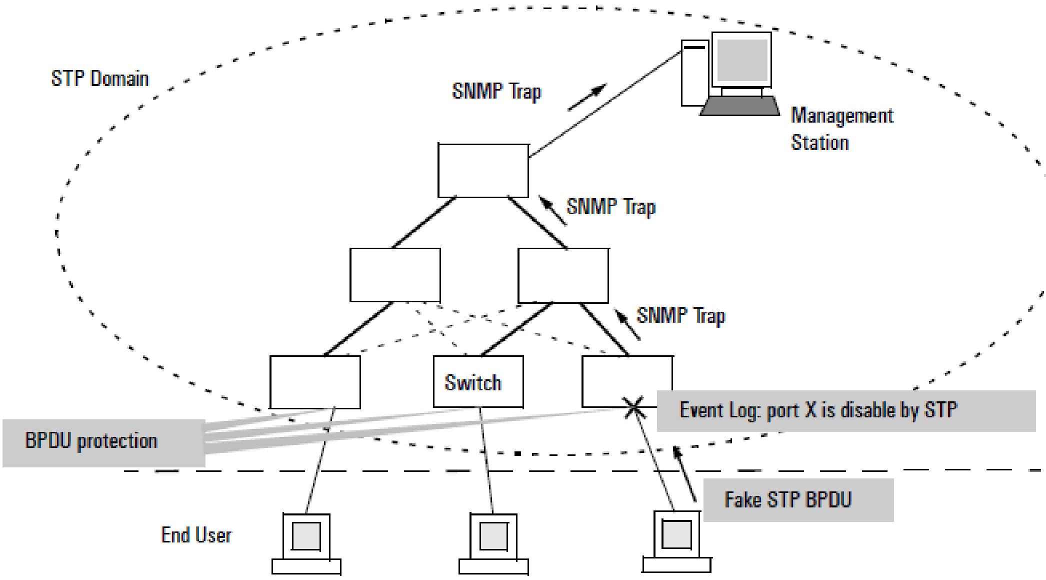

The following commands allow you to configure BPDU protection on VLANs for which the port is a member.

Syntax:

```
no spanning-tree port-list bpdu-protection
```

Enables/disables the BPDU protection feature on a port.

Default: Disabled.

Syntax:

```
no spanning-tree port-list bpdu-protection-timeout timeout
```

Configures the duration of time when protected ports receiving unauthorized BPDUs will remain disabled. The default value of 0 (zero) sets an infinite timeout (that is, ports that are disabled by `bpdu-protection` are not, by default, re-enabled automatically).

Default: 0

Range: 0 - 65535 seconds

Syntax:

```
no spanning-tree trap errant-bpdu
```

Enables/disables the sending of errant BPDU traps.

CAUTION: This command should only be used to guard edge ports that are not expected to participate in STP operations. Once BPDU protection is enabled, it will disable the port as soon as any BPDU packet is received on that interface.

Viewing BPDU protection status

Syntax:

`show spanning-tree bpdu-protection [`*`port-list`*`]`

Displays a summary listing of ports with BPDU protection enabled. To display detailed per-port status information, enter the specific port number(s). BPDU protected ports are displayed as separate entries of the spanning tree category within the configuration file.

Viewing BPDU protection status for specific ports

```
switch# show spanning-tree bpdu-protection 23-24

 Status and Counters - STP BPDU Protection Information

 BPDU Protection Timeout (sec) : 0
 BPDU Protected Ports : 23-24


  Port   Type      Protection State             Errant BPDUs
  ------ --------- ---------- ----------------- ------------
  23     100/1000T Yes        Bpdu Error        1
  24     100/1000T Yes                          0
```

Re-enabling a port blocked by BPDU protection

Ports disabled by BPDU Protection remain disabled unless BPDU Protection is removed from the switch or by configuring a nonzero BPDU protection timeout. For example, if you want to re-enable protected ports 60 seconds after receiving a BPDU, you would use this command:

```
switch(config)# spanning-tree bpdu-protection-timeout 60
```

Enabling and disabling BPDU protection

Syntax:

`no spanning-tree `*`port-list`*` bpdu-protection`

Enables or disables BPDU protection on specified port(s).

Syntax:

`no spanning-tree `*`port-list`*` bpdu-protection-timeout `*`timeout`*

Configures the duration in seconds when protected ports receiving unauthorized BPDUs will remain disabled. The default value of 0 (zero) sets an infinite timeout (that is, ports that are disabled by `bpdu-protection` are not, by default, re-enabled automatically).

Range: 0-65535 seconds

Default: 0

Syntax:

`no spanning-tree trap errant-bpdu`

Enables or disables the sending of errant BPDU traps.

CAUTION: This command should only be used to guard edge ports that are not expected to participate in STP operations. Once BPDU protection is enabled, it will disable the port as soon as any BPDU packet is received on that interface.

Configuring BPDU protection

To configure BPDU protection on ports 1 to 10 with SNMP traps enabled, enter:

```
switch(config)# spanning-tree 1-10 bpdu protection
switch(config)# spanning-tree trap errant-bpdu
```

The following steps will then be set in progress:

1. When an STP BPDU packet is received on ports 1-10, STP treats it as an unauthorized transmission attempt and shuts down the port that the BPDU came in on.
2. An event message is logged and an SNMP notification trap is generated.
3. The port remains disabled until re-enabled manually by a network administrator using the `interface` *port-list* `enable` command.

NOTE: To re-enable the BPDU-protected ports automatically, configure a timeout period using the `spanning-tree bpdu-protection-timeout` command.

Overview of MSTP BPDU throttling

When an STP enabled switch is hit by an MSTP BPDU storm, the CPU usage rises and the manageability of the switch goes down. In the following figure, the switch is connected to a HUB where there is a loop. The switch generates a single MSTP BPDU, which goes through the loop in the HUB and results in a BPDU storm eventually. Since all STP packets are taken to the CPU for processing, CPU usage goes high and the switch response slows down. The switch can become unmanageable as a result of this BPDU storm. BPDU throttling will take care of BPDU storms automatically through rate-limiting.

Figure 21: *MSTP BPDU path*

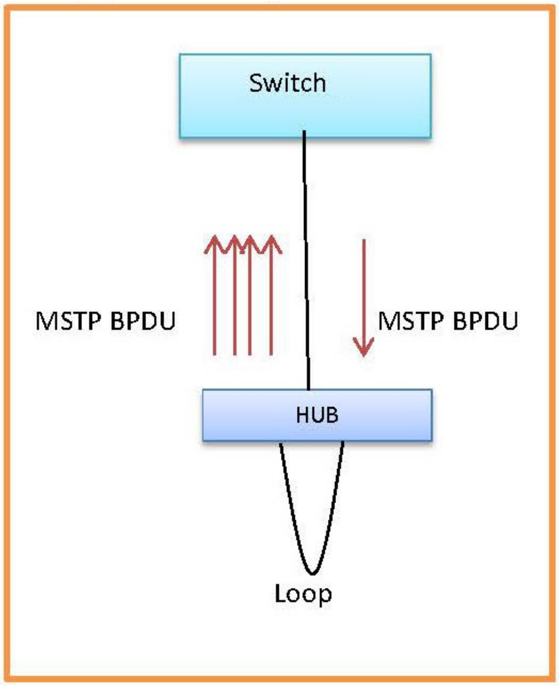

BPDU throttling is enabled when the spanning tree in MSTP mode is enabled. When spanning tree is enabled, all modules and members are assigned corresponding throttle values from the configuration. The default throttle value is 256.

An option is also provided to enabling/disabling BPDU throttling. This option is enabled by default if the switch does not support "V1 modules". The spanning tree is enabled in MSTP mode by default.

Configuring MSTP BPDU throttling

The CLI allows you to configure MSTP BPDU throttling.

Configuring MSTP BPDU throttling

Syntax

```
no spanning-tree bpdu-throttle [Throttle-Value]
```

Configures BPDU throttling on a device. BPDU throttling limits the number of BPDUs that are sent to the switch's CPU. The result prevents high CPU utilization on the switch when the network undergoes a broadcast storm or loop. The BPDU throttle value is in packets per second (pps). The valid BPDU throttle values are 64, 128, and 256 pps. The default throttle value is 256 pps.

Show MSTP BPDU configuration

The CLI allows you to show MSTP BPDU throttling configurations.

Syntax

```
show spanning-tree bpdu-throttle
```

Displays the configured throttle value.

Example

```
Show spanning-tree bpdu-throttle
BPDU Throttling State  : Enabled
BPDU Throttle value    : 256
```

Show running configuration

Syntax

```
show running configuration
```

Show running configuration will display any one of the following lines based on the configuration.

```
no spanning-tree bpdu throttle
spanning-tree bpdu throttle 128
spanning-tree bpdu throttle 64
```

PVST

PVST stands for Per-VLAN Spanning Tree. It allows for the creation of a spanning tree for each VLAN.

PVST protection and filtering

NOTE:

These options are available for switches that support the MSTP protocol only. They are not supported for switches running RSTP.

PVST protection

If a switch in the core of a network receives Per Vlan Spanning Tree (PVST) BPDUs and forwards the unrecognized PVST BPDUs on to MSTP-only switches, those switches then disconnect themselves from the network. This can create instability in the network infrastructure.

When the PVST protection feature is enabled on a port and a PVST BPDU is received on that port, the interface on which the PVST BPDU arrived is shut down, which isolates the sending switch from the rest of the network. An event message is logged and an SNMP notification trap is generated. The errant BPDU counter `SwitchStpPortErrantBpduCounter` is incremented. The PVST protection feature is enabled per-port.

Figure 22: *PVST switch being isolated after sending a PVST BPDU*

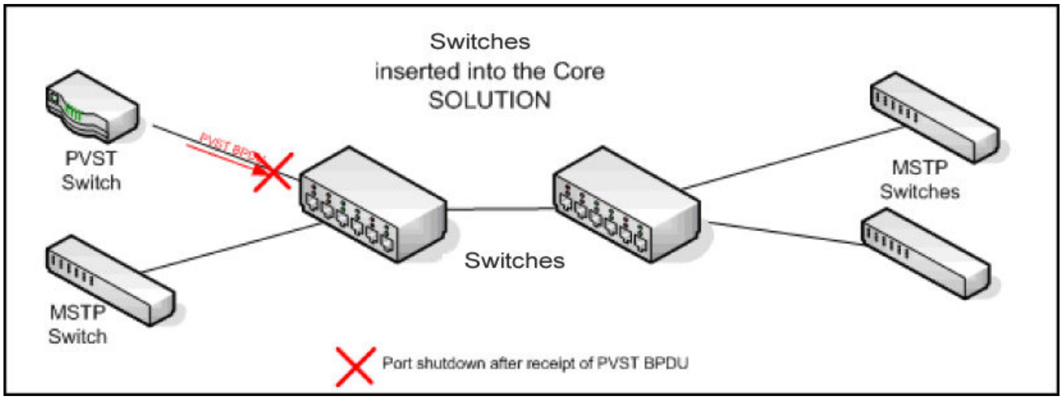

NOTE: This is similar to the BPDU Guard feature where BPDU protection is applied to edge ports connected to end user devices that do not run STP. If STP BPDU packets are received on a protected port, the feature will disable that port and alert the network manager via an SNMP trap.

PVST filtering

If you configure a port for PVST filtering instead of PVST protection, the port remains in operation but traps are still generated and the BPDU counter `SwitchStpPortErrantBpduCounter` is incremented.

CAUTION: Enabling the PVST filter feature allows the port to continuously forward packets without spanning tree intervention, which could result in loop formation. If this occurs, disable the port and then reconfigure it with these commands:

```
no spanning-tree port-list bpdu-filter
```

```
no spanning-tree port-list pvst-filter
```

Enabling and disabling PVST protection on ports

Syntax:

```
no spanning-tree port-list pvst-protection
```

Enables or disables PVST protection on the port(s) specified. The command indicates which ports are not expected to receive any PVST BPDUs.

Default: Disabled on all ports.

Enabling PVST protection

To enable the PVST protection feature on ports 4 through 8, enter:

```
switch(config)# spanning-tree 4-8 pvst-protection
```

To disable the PVST protection feature on a port, for example, port 4, enter:

```
switch(config)# no spanning-tree 4 pvst-protection
```

Enabling and disabling PVST filters on ports

Syntax:

```
no spanning-tree port-list pvst-filter
```

Enables or disables PVST filters on the port(s) specified. The command indicates which ports are not expected to receive any PVST BPDUs.

Default: Disabled on all ports

Enabling PVST filtering on a port

```
switch(config)# spanning-tree 8 pvst-filter

Warning: The BPDU filter allows the port to go into a continuousforwarding mode
and spanning-tree will not interfere, even if the port would cause a loop to form
in the network topology.
If you suddenly experience high traffic load, disable the port and reconfigure the
```

```
BPDU filter with the CLI command(s):
            "no spanning-tree PORT_LIST bpdu-filter"
            "no spanning-tree PORT_LIST pvst-filter"
```

Re-enabling a port manually

Syntax:

```
no spanning-tree bpdu-protection-timeout timeout
```

Configures the duration of time protected ports remain disabled. The default value of `0` sets an infinite timeout, so ports that are disabled are not re-enabled automatically.

NOTE: This is a GLOBAL command.

Range: 0 - 65535 seconds

Default: 0

You can also set the timeout in the MIB with this MIB object: `hpSwitchStpBpduProtectionTimeout`

It is also possible to use the following automatic re-enable timer command:

```
switch(config)# spanning-tree bpdu-protection-timeout 120
```

Viewing ports configured with PVST protection and filtering

Viewing all ports with PVST protection enabled

```
switch(config)# show spanning-tree pvst-protection

Status and Counters - PVST Port(s) BPDU Protection Information

BPDU Protection Timeout (sec) : 0
PVST Protected Ports : 5-6
```

Viewing all ports with PVST filtering enabled

```
switch(config)# show spanning-tree pvst-filter
Status and Counters - PVST Port(s) BPDU Filter Information
PVST Filtered Ports : 8
```

Listing ports to see which have PVST protection or filtering enabled

Syntax:

```
show spanning-tree <port-list> detail
```

Viewing if PVST protection is enabled (Yes)

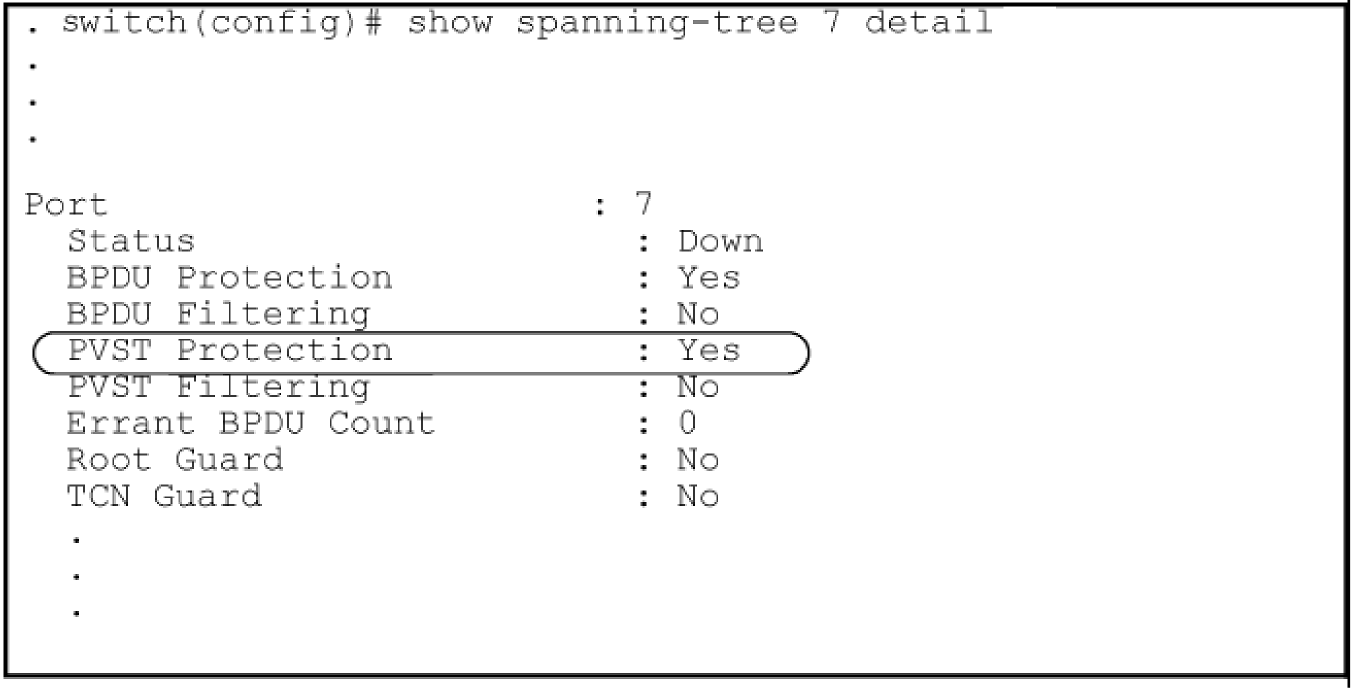

```
. switch(config)# show spanning-tree 7 detail
.
.
.

Port                          : 7
  Status                        : Down
  BPDU Protection               : Yes
  BPDU Filtering                : No
  PVST Protection               : Yes
  PVST Filtering                : No
  Errant BPDU Count             : 0
  Root Guard                    : No
  TCN Guard                     : No
  .
  .
  .
```

In cases where spanning tree cannot be used to prevent loops at the edge of the network, loop protection may provide a suitable alternative. Loop protection operates in two modes:

Untagged

The default mode. This mode can be used to find loops in untagged downlinks.

Tagged VLAN

Finds loops on tagged VLANs. This mode can be used to detect loops in tagged-only uplinks where STP cannot be enabled.

The cases where loop protection might be chosen ahead of spanning tree to detect and prevent loops are as follows:

On ports with client authentication

When spanning tree is enabled on a switch that use 802.1X, Web authentication, and MAC authentication, loops may go undetected. For example, spanning tree packets that are looped back to an edge port will not be processed because they have a different broadcast/multicast MAC address from the client-authenticated MAC address. To ensure that client-authenticated edge ports get blocked when loops occur, you should enable loop protection on those ports.

On ports connected to unmanaged devices

Spanning tree cannot detect the formation of loops where there is an unmanaged device on the network that does not process spanning tree packets and simply drops them. Loop protection has no such limitation, and can be used to prevent loops on unmanaged switches.

Loop protection enabled in preference to STP

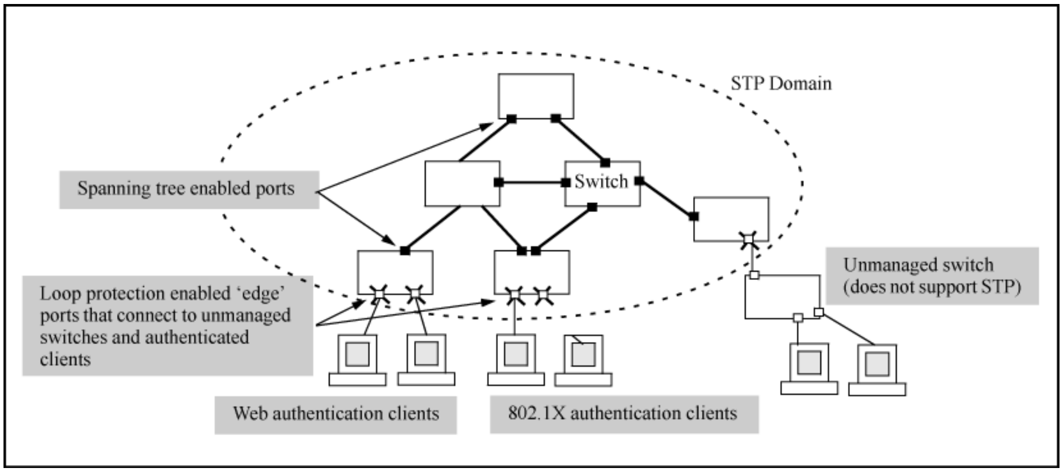
STP Domain
Spanning tree enabled ports
Switch
Loop protection enabled 'edge' ports that connect to unmanaged switches and authenticated clients
Unmanaged switch (does not support STP)
Web authentication clients
802.1X authentication clients

Configuring loop protection

Loop protection provides protection against loops by transmitting loop protocol packets out of ports on which loop protection has been enabled. When the switch sends out a loop protocol packet and then receives the same packet on a port that has a `receiver-action` of `send-disable` configured, it shuts down the port from which the packet was sent.

Syntax:

```
no loop-protect port-list [[receiver-action [[send-disable] | [no-disable]]] |
[transmit-interval 1-10] | [disable-timer 0-604800] | [trap loop-detected]] [mode]
[[port] | [vlan]] [vlan vid-list]
```

Description

Configures per-port loop protection on the switch.

Parameters

`receiver-action send-disable | no-disable`

Sets the action to be taken when a loop is detected on the specified ports. The port that receives the loop protection packet determines what action is taken. If `send-disable` is configured, the port that transmitted the packet is disabled. If `no-disable` is configured, the port is not disabled when a loop is detected.

NOTE: The port will not transmit loop protection packets unless it is a member of an untagged VLAN. If a port is only a member of tagged VLANs, the loop protection packets are not transmitted.

Default: `send-disable`

`trap loop-detected`

Configures loop protection traps for SNMP indicating when a loop has been detected on a port.

`disable-timer` ***`0-604800`***

Configures how long (in seconds) a port is disabled when a loop has been detected. A value of zero disables the auto re-enable function.

Default: Timer is disabled

`transmit-interval` ***`1-10`***

Configures the time in seconds between the transmission of loop protection packets.

Default: 5 seconds

`{mode port | vlan}`

Configures loop protection in port or VLAN mode.

`vlan` ***`vid-list`***

Configures the VLANs on which loop-protect is enabled. Maximum number of loop-protected VLANs is 32.

Enabling loop protection in port mode

Follow these steps.

Procedure

1. Configure port mode with this command:

```
switch(config)# loop-protect mode port
```

2. Enter the `loop-protect` command and specify the ports on which loop protection should be enabled. For example:

```
switch(config)# loop-protect 1-2
```

3. Optionally specify `receiver-action` of `send-disable` to shut down the port in the event of a loop. For example:

```
switch(config)# loop-protect 1-2 receiver-action send-disable
```

Enabling loop protection in VLAN mode

VLANs can be configured for loop protection only when operating in VLAN mode. When `loop-protect` is enabled for a VLAN and a `loop-protect` enabled interface is a member of that VLAN, loop protect packets are sent on that VLAN to detect loops.

To enable loop protection in VLAN mode:

Procedure

1. Configure VLAN mode with the command:

```
switch(config)# loop-protect mode vlan
```

2. Enter the `loop-protect` command and specify the VLANs on which loop protection should be enabled. For example:

```
switch(config)# loop-protect vlan 20,30
```

Changing modes for loop protection

When changing from VLAN mode to port mode, you are prompted with the message shown below. The VLANs will no longer be configured for loop protection.

Changing modes for loop protection

```
switch(config)# loop-protect mode port
Any Loop Protect enabled VLAN will be deleted. Do you want to continue [Y/N]?
N
```

Viewing loop protection status in port mode

Syntax:

`show loop-protect`*`port-list`*

Displays the loop protection status for ports. If no ports are specified, the information is displayed only for the ports that have loop protection enabled.

Viewing loop protection information for port mode

```
switch(config)# show loop-protect 1-2

 Status and Counters - Loop Protection Information

 Transmit Interval (sec)     : 5
 Port Disable Timer (sec)    : 5
 Loop Detected Trap          : Enabled
 Loop Protect Mode           : Port
 Loop Protect Enabled VLANs :
     Loop    Loop     Detected  Loop     Time Since  Rx          Port
Port Protect Detected on VLAN   Count    Last Loop   Action      Status
---- ------- -------- --------- -------- ----------- ---------- ---------
1    Yes     Yes        NA      1        5s         send-disable Down
2    Yes     No         NA      0                   send-disable Up
```

Viewing loop protection status in VLAN mode

Syntax:

```
show loop-protect port-list
```

Displays the loop protection status for VLANs. If no ports are specified, the information is displayed only for the ports that have loop protection enabled.

Viewing loop protection information for VLAN mode

```
switch(config)# show loop-protect 1-2

 Status and Counters - Loop Protection Information

 Transmit Interval (sec)   : 5
 Port Disable Timer (sec)  : 5
 Loop Detected Trap        : Enabled
 Loop Protect Mode         : Vlan
 Loop Protect Enabled VLANs : 20,30


      Loop   Loop     Detected  Loop     Time Since  Rx          Port
Port Protect Detected on VLAN   Count    Last Loop   Action      Status
---- ------- -------- --------- -------- ----------- ---------- ---------
1    Yes     Yes        20      1        45s        send-disable Down
2    Yes     No                 0                   send-disable Up
```

STP loop guard

Spanning Tree (STP) is used to ensure a loop-free topology over the LAN. Occasionally a hardware or software failure can cause STP to fail, creating STP/forwarding loops that can cause network failures where unidirectional links are used. The non-designated port transitions in a faulty manner because the port is no longer receiving STP BPDUs.

STP Loop Guard causes the non-designated port to go into the STP loop inconsistent state instead of the forwarding state. In the loop-inconsistent state, the port prevents data traffic and BPDU transmission through the link, therefore avoiding the loop creation. When BPDUs again are received on the inconsistent port, it resumes normal STP operation automatically. STP loop guard is best applied on blocking or forwarding ports.

Figure 23: *Loop creation with transmission failure*

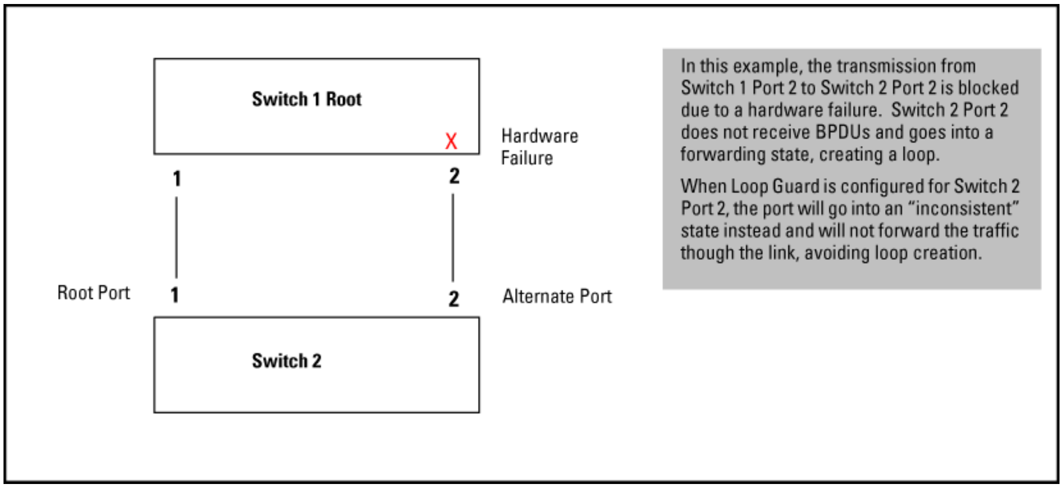

In this example, the transmission from Switch 1 Port 2 to Switch 2 Port 2 is blocked due to a hardware failure. Switch 2 Port 2 does not receive BPDUs and goes into a forwarding state, creating a loop.

When Loop Guard is configured for Switch 2 Port 2, the port will go into an "inconsistent" state instead and will not forward the traffic though the link, avoiding loop creation.

Syntax:

`spanning-tree` *`port-list`* `loop-guard`

`no spanning-tree` *`port-list`* `loop-guard`

Enables STP loop guard on a particular port or ports. The `no` form of the command disables STP loop guard.

Default: Disabled.

Enabling spanning tree loop guard on Port 2 and Viewing the port's status

```
switch(config)# spanning-tree 2 loop-guard
```

```
switch(config)# show spanning-tree

 Multiple Spanning Tree (MST) Information

  STP Enabled   : Yes
  Force Version : MSTP-operation
  IST Mapped VLANs : 1-4094
  Switch MAC Address : 0024a8-d13a40
  Switch Priority    : 32768
  Max Age  : 20
  Max Hops : 20
  Forward Delay : 15

  Topology Change Count  : 1
  Time Since Last Change : 20 mins

  CST Root MAC Address : 001083-847000
  CST Root Priority    : 0
  CST Root Path Cost   : 60000
  CST Root Port        : 1

  IST Regional Root MAC Address : 0024a8-d13a40
  IST Regional Root Priority    : 32768
  IST Regional Root Path Cost   : 0
  IST Remaining Hops            : 20

  Root Guard Ports     :
```

```
  Loop Guard Ports     : 2
  TCN Guard Ports      :
  BPDU Protected Ports :
  BPDU Filtered Ports  :
  PVST Protected Ports :
  PVST Filtered Ports  :

                 |           Prio              | Designated     Hello
Port   Type      | Cost      rity State        | Bridge         Time PtP Edge
------ --------- + --------- ---- ------------ + -------------- ---- --- ----
1      100/1000T | 20000     128  Forwarding   | 001871-cdea00  2    Yes No
2      100/1000T | Auto      128  Inconsistent |
3      100/1000T | Auto      128  Disabled     |
4      100/1000T | Auto      128  Disabled     |
5      100/1000T | Auto      128  Disabled     |
6      100/1000T | Auto      128  Disabled     |
7      100/1000T | Auto      128  Disabled     |
8      100/1000T | Auto      128  Disabled     |
```

Viewing summary spanning tree configuration information

```
switch(config)# show spanning-tree config

 Multiple Spanning Tree (MST) Configuration Information

  STP Enabled [No] : Yes
  Force Version [MSTP-operation] : MSTP-operation
  Default Path Costs [802.1t] : 802.1t
  MST Configuration Name : 0024a8d13a40
  MST Configuration Revision : 0       Switch Priority : 32768
  Forward Delay [15] : 15              Hello Time [2] : 2
  Max Age [20] : 20                    Max Hops [20] : 20

               | Path      Prio Admin Auto Admin Hello  Root  Loop  TCN   BPDU
Port Type      | Cost      rity Edge  Edge PtP   Time   Guard Guard Guard Flt
---- --------- + --------- ---- ----- ---- ----- ------ ----- ----- ----- ---
1    100/1000T | Auto      128  No    Yes  True  Global No    No    No    No
2    100/1000T | Auto      128  No    Yes  True  Global No    Yes   No    No
3    100/1000T | Auto      128  No    Yes  True  Global No    No    No    No
4    100/1000T | Auto      128  No    Yes  True  Global No    No    No    No
5    100/1000T | Auto      128  No    Yes  True  Global No    No    No    No
6    100/1000T | Auto      128  No    Yes  True  Global No    No    No    No
.
.
.
```

Viewing detailed spanning tree configuration information

```
switch(config)# show spanning-tree detail

 Status and Counters - CST Port(s) Detailed Information
  Port                     : 1
  Status                   : Up
.
.
.

  Port                     : 2
  Status                   : Up
  BPDU Protection          : No
  BPDU Filtering           : No
  PVST Protection          : No
  PVST Filtering           : No
  Errant BPDU Count        : 0
  Root Guard               : No
  Loop Guard               : Yes
  TCN Guard                : No
  MST Region Boundary      : Yes
  External Path Cost       : 20000
  External Root Path Cost  : 40000
  Administrative Hello Time: Global
  Operational Hello Time   : 2
  AdminEdgePort            : No
  Auto Edge Port           : Yes
  OperEdgePort             : No
  AdminPointToPointMAC     : True
  OperPointToPointMAC      : Yes
  Aged BPDUs Count         : 0
  Loop-back BPDUs Count    : 0
  TC ACK Flag Transmitted  : 0
  TC ACK Flag Received     : 1

 MST        MST        CFG        CFG        TCN        TCN
 BPDUs Tx   BPDUs Rx   BPDUs Tx   BPDUs Rx   BPDUs Tx   BPDUs Rx
 ---------- ---------- ---------- ---------- ---------- ----------
 3          0          24354      1682       0          13
```

Viewing spanning tree configuration information for a single port

```
switch(config)# show spanning-tree 2

 Multiple Spanning Tree (MST) Information

  STP Enabled   : Yes
  Force Version : MSTP-operation
  IST Mapped VLANs : 1-4094
  Switch MAC Address : 0024a8-d13a40
  Switch Priority   : 32768
  Max Age : 20
  Max Hops : 20
  Forward Delay : 15

  Topology Change Count  : 1
  Time Since Last Change : 58 mins

  CST Root MAC Address : 001083-847000
  CST Root Priority   : 0
  CST Root Path Cost  : 60000
  CST Root Port       : 1

  IST Regional Root MAC Address : 0024a8-d13a40
  IST Regional Root Priority   : 32768
  IST Regional Root Path Cost  : 0
  IST Remaining Hops           : 20

  Root Guard Ports     :
  Loop Guard Ports     : 2
  TCN Guard Ports      :
  BPDU Protected Ports :
  BPDU Filtered Ports  :
  PVST Protected Ports :
  PVST Filtered Ports  :

                 |           Prio               | Designated    Hello
Port   Type      | Cost      rity State         | Bridge        Time PtP Edge
```

```
------ --------- + --------- ---- ------------ + ------------- ---- --- ----
2      100/1000T | Auto      128  Inconsistent |
```

Operating notes

- The `receiver-action` option can be configured on a per-port basis and can only be enabled after loop protection has been enabled on the port. All other configuration options (disable-timer, trap loop-detected, and transmit interval) are global.
- The `trap` option refers to an SNMP trap.
- Regardless of how the `receiver-action` and `trap` options are configured, all detected loops will be logged in the switch's event log.
- The `no loop-protect port` command will not remove a receive-action configuration line from the running configuration unless this option is set to `receive-action send-disable`.
- If `loop-protect` is enabled in port mode, it cannot also be enabled in VLAN mode, and vice-versa.

Chapter 8
Quality of Service (QoS): Managing bandwidth effectively

Introduction to Quality of Service (QoS)

A Quality of Service (QoS) **network policy** refers to the network-wide controls available to:

- Ensure uniform and efficient traffic-handling throughout your network, while keeping the most important traffic moving at an acceptable speed, regardless of current bandwidth usage.
- Exercise control over the priority settings of inbound traffic arriving in and travelling through your network.

Adding bandwidth can be a good idea, but is not always feasible and does not completely eliminate the potential for network congestion. There will always be points in the network where multiple traffic streams merge or where network links change speed and capacity. The impact and number of these congestion points will increase over time as more applications and devices are added to the network.

When network congestion occurs, it is important to move traffic on the basis of relative importance. However, without QoS prioritization, less important traffic consumes network bandwidth and slows down or halts the delivery of more important traffic. Without QoS, most traffic received by the switch is forwarded with the same priority it had upon entering the switch. In many cases, such traffic is normal priority and competes for bandwidth with all other normal-priority traffic, regardless of its relative importance to your organization's mission.

Using QoS to classify and prioritize network traffic

QoS is used to classify and prioritize traffic throughout a network. QoS enables you to establish an end-to-end traffic-priority policy to improve the control and throughput of important data. You can manage available bandwidth so that the most important traffic goes first. For example, you can use QoS to:

- Upgrade or downgrade traffic from various servers.
- Control the priority of traffic from dedicated VLANs or applications.

- Change the priorities of traffic from various segments of your network as your business needs change.
- Set priority policies in edge switches in your network to enable traffic-handling rules across the network.

Figure 24: *802.1p priority based on CoS (Class-of-Service) types and use of VLAN tags*

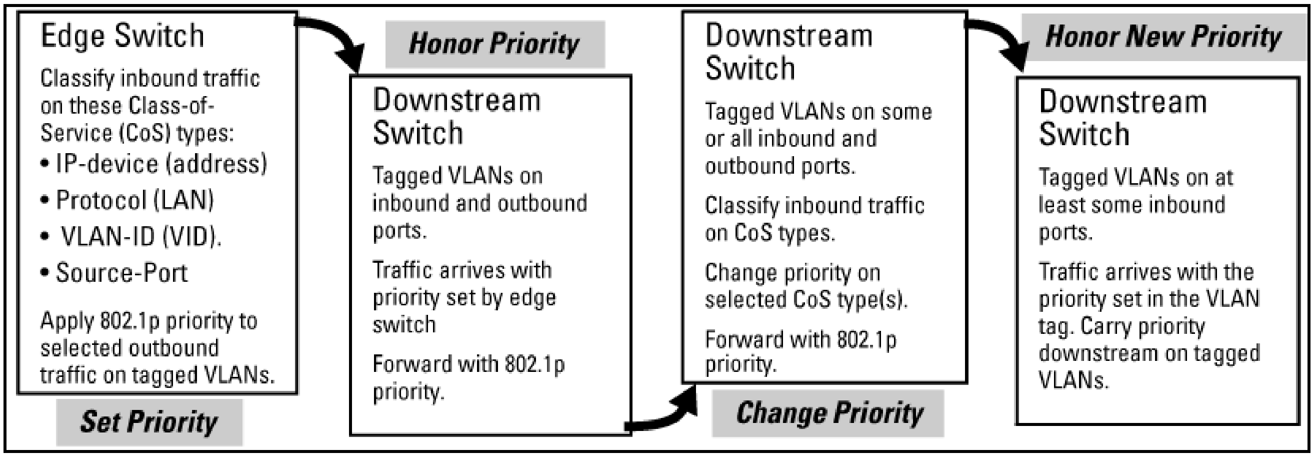
Edge Switch
Classify inbound traffic on these Class-of-Service (CoS) types:
• IP-device (address)
• Protocol (LAN)
• VLAN-ID (VID).
• Source-Port
Apply 802.1p priority to selected outbound traffic on tagged VLANs.
Set Priority
Honor Priority
Downstream Switch
Tagged VLANs on inbound and outbound ports.
Traffic arrives with priority set by edge switch
Forward with 802.1p priority.
Downstream Switch
Tagged VLANs on some or all inbound and outbound ports.
Classify inbound traffic on CoS types.
Change priority on selected CoS type(s).
Forward with 802.1p priority.
Change Priority
Honor New Priority
Downstream Switch
Tagged VLANs on at least some inbound ports.
Traffic arrives with the priority set in the VLAN tag. Carry priority downstream on tagged VLANs.

Figure 25: *Application of Differentiated Services Codepoint (DSCP) policies*

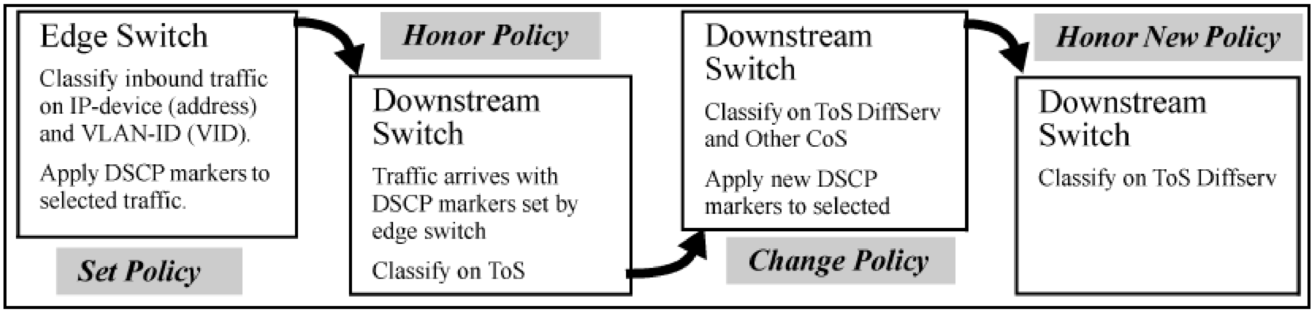
Edge Switch
Classify inbound traffic on IP-device (address) and VLAN-ID (VID).
Apply DSCP markers to selected traffic.
Set Policy
Honor Policy
Downstream Switch
Traffic arrives with DSCP markers set by edge switch
Classify on ToS
Downstream Switch
Classify on ToS DiffServ and Other CoS
Apply new DSCP markers to selected
Change Policy
Honor New Policy
Downstream Switch
Classify on ToS Diffserv

Applying QoS to inbound traffic at the network edge

At the edge switch, QoS classifies certain traffic types and in some cases applies a DSCP policy. At the next hop (downstream switch) QoS honors the policies established at the edge switch. Further downstream, another switch may reclassify some traffic by applying new policies, and yet other downstream switches can be configured to honor the new policies.

Preserving QoS in outbound traffic in a VLAN

QoS is implemented in the form of rules or policies that are configured on the switch. Although you can use QoS to prioritize traffic only while it moves through the switch, you derive the maximum benefit by using QoS in an 802.1Q VLAN environment (with 802.1p priority tags) or in an untagged VLAN environment (with DSCP policies in which QoS sets priorities that downstream devices can support without reclassifying the traffic).

Using QoS to optimize existing network resources

By prioritizing traffic, QoS supports traffic growth on the network while optimizing the use of existing resources—and delaying the need for further investments in equipment and services. QoS enables you to:

- Specify which traffic has higher or lower priority, regardless of current network bandwidth, or the relative priority setting of the traffic when it is received on the switch.
- Change (upgrade or downgrade) the priority of outbound traffic.
- Override "illegal" packet priorities set by upstream devices or applications that use 802.1Q VLAN tagging with 802.1p priority tags.

Overview of QoS settings

QoS settings operate on two levels:

- **Controlling the priority of outbound packets moving through the switch:** Configuring a new 802.1p priority value allows you to set the outbound priority queue to which a packet is sent. For example, you can configure an 802.1p priority of 0 through 7 for an outbound packet. When the packet is sent to a port, the QoS priority determines the outbound queue to which the packet is assigned as shown in the following table:

  Table 16: *802.1p priority settings and outbound queue assignment*

802.1p priority setting	Outbound port 8 queues	Outbound port 4 queues	Outbound port 2 queues
1	1	1	1
2	2		
0	3	2	
3	4		
4	5	3	2
5	6		
6	7	4	
7	8		

  (In an 802.1Q VLAN environment with VLAN-tagged ports, if QoS is **not** configured on the switch, but is configured on an upstream device, the priorities carried in the packets determine the forwarding queues in the switch.)

- **Configuring a priority for outbound packets and a service (priority) policy for use by downstream devices:**
  - **DSCP Policy:** This feature enables you to set a priority policy in outbound IP packets. (You can configure downstream devices to read and use this policy.) This method is not dependent on VLAN-tagged ports to carry priority policy to downstream devices, and can:
    - Change the codepoint (the upper 6 bits) in the ToS byte.
    - Set a new 802.1p priority for the packet.

    (Setting DSCP policies requires IPv4 inbound packets.)
  - **802.1p priority rules:** An outbound, VLAN-tagged packet carries an 802.1p priority setting that was configured (or preserved) in the switch. This priority setting ranges from 0 to 7, and can be used by downstream devices having up to eight outbound port queues. Thus, if packets within the switch move at the four priority levels shown in the table above, they still can carry an 802.1p priority that can be used by downstream devices having more or less than the four priority levels in the switches

covered in this guide. Also, if the packet enters the switch with an 802.1p priority setting, QoS can override this setting if configured with an 802.1p priority rule to do so.

NOTE: If your network uses only one VLAN (and therefore does not require VLAN-tagged ports), you can still preserve 802.1p priority settings in your traffic by configuring the ports as tagged VLAN members on the links between devices you want to honor traffic priorities.

Rule and policy limits: A large number of 802.1p priority rules and/or DSCP policies are allowed in any combination. For example, for the 2540 switch 6000 are allowed.

You can configure a QoS priority of 0 through 7 for an outbound packet. When the packet is then sent to a port, the QoS priority determines which outbound queue the packet uses:

Table 17: *QoS priority settings and operation*

802.1p priority setting	Outbound port 8 queues	Outbound port 4 queues	Outbound port 2 queues
1	1	1	1
2	2		
0	3	2	
3	4		
4	5	3	2
5	6		
6	7	4	
7	8		

If a packet is not in a VLAN-tagged port environment, then the QoS settings in the table above control only to which outbound queue the packet goes. Without VLAN tagging, no 802.1p priority is added to the packet for downstream device use. But if the packet is in a VLAN-tagged environment, then the above setting is also added to the packet as an 802.1p priority for use by downstream devices and applications (shown in the table below). In either case, an IP packet can also carry a priority policy to downstream devices by using DSCP-marking in the ToS byte.

Table 18: *Mapping switch QoS priority settings to device queues*

Priority setting	Outbound port queues in the switch	802.1p priority setting added to tagged VLAN packets exiting the switch	Queue assignment in downstream devices with:		
			8 queues	3 queues	2 queues
1	Queue 1	1 (low priority)	Queue 1	Queue 1	Queue 1
2		2	Queue 2		
0	Queue 2	0 (normal priority)	Queue 3	Queue 2	

Table Continued

Priority setting	Outbound port queues in the switch	802.1p priority setting added to tagged VLAN packets exiting the switch	Queue assignment in downstream devices with:		
			8 queues	3 queues	2 queues
3		3	Queue 4		
4	Queue 3	4 (medium priority)	Queue 5	Queue 3	Queue 2
5		5	Queue 6		
6	Queue 4	6 (high priority)	Queue 7		
7		7	Queue 8		

Classifiers for prioritizing outbound packets

NOTE:

Regarding using multiple criteria: Hewlett Packard Enterprise recommends that you configure a minimum number of the available QoS classifiers for prioritizing any given packet type. Increasing the number of active classifier options for a packet type increases the complexity of the possible outcomes and consumes switch resources.

Packet classifiers and evaluation order

The switches covered in this guide provide six types of globally-configured QoS classifiers (match criteria) to select packets for QoS traffic marking.

The switches covered in this guide provide six QoS classifiers (packet criteria) you can use to configure QoS priority.

Table 19: *Classifier search order and precedence*

Search order	Precedence	Global QoS classifier
1	1 (highest)	UDP/TCP application type (port)
2	2	Device priority (destination or source IP address)
3	3	IP type of service (ToS): precedence and DSCP bit sets (IP packets only)
4	4	IP protocol (IP, IPX, ARP, AppleTalk, SNA, and NetBeui)
5	5	VLAN ID

Table Continued

Search order	Precedence	Global QoS classifier
6	6	Incoming source-port on the switch
Default	7 (lowest)	The incoming 802.1p priority (present in tagged VLAN environments) is preserved if no global QoS classifier with a higher precedence matches.

The switch uses the lowest-to-highest search order shown in the table to identify the highest-precedence classifier to apply to any given packet. If there is only one configured type that matches a given packet, then the switch applies the QoS policy specified in that type. If multiple configured types match a given packet, the switch applies each one in turn to the packet and concludes with the QoS policy for the highest-precedence type.

- If the highest-precedence classifier is configured to apply a DSCP policy, then both the DSCP in the packet and the 802.1p priority applied to the packet can be changed.
- If the highest-precedence classifier is configured to apply an 802.1p priority rule, then only the 802.1 priority in the final QoS match for the packet can be changed.

NOTE: Intermixing lower-precedence types configured with DSCP policies and higher-precedence types configured with 802.1p priority rules, is **not recommended**, as this can result in a packet with an 802.1p priority assigned by one type and a DSCP policy by another type. This is because the search order would allow a lower-precedence type configured with a DSCP policy to change both the DSCP and the 802.1p setting in a packet, and then would allow a subsequent, higher-precedence type configured with an 802.1p priority rule to change only the 802.1p setting.

To avoid this problem, a DSCP policy option should be applied only on the highest-precedence type in use on the switch, or apply to all QoS types in use on the switch.

Preparation for configuring QoS

Preserving 802.1p priority

QoS operates in VLAN-tagged and VLAN-untagged environments. If your network does not use multiple VLANs, you can still implement the 802.1Q VLAN capability for packets to carry their 802.1p priority to the next downstream device. To do so, configure ports as VLAN-tagged members on the links between switches and routers in your network infrastructure.

Table 20: *Summary of QoS capabilities*

Outbound packet options	Port membership in VLANs	
	Tagged	Untagged
Control port queue priority for packet types	Yes	Yes
Carry 802.1p priority assignment to next downstream device	Yes	No
Carry DSCP policy to downstream devices. The policy includes: • Assigning a ToS Codepoint • Assigning an 802.1p Priority to the Codepoint	Yes	Yes[2]

Steps for configuring QoS on the switch

Procedure

1. Determine the QoS policy you want to implement. This includes analyzing the types of traffic flowing through your network and identifying one or more traffic types to prioritize. In order of QoS precedence, these are:

   a. UDP/TCP applications

   b. Device Priority—destination or source IP address (Note that destination has precedence over source. See the table below.

   c. IP ToS Precedence Bits (Leftmost three bits in the ToS field of IP packets)

   d. IP ToS Differentiated Service bits (Leftmost 6 bits in the ToS field of IP packets)

   e. Layer 3 Protocol Priority

   f. VLAN Priority (requires at least one tagged VLAN on the network)

   g. Source-Port

   h. Incoming 802.1p Priority (requires at least one tagged VLAN on the network). In a tagged VLAN environment, the incoming 802.1p priority is used as the default QoS classifier if no QoS classifier with a higher precedence matches

2. Select the QoS option you want to use. The following table lists the traffic types (QoS classifiers) and the QoS options you can use for prioritizing or setting a policy on these traffic types:

Table 21: *QoS marking supported by QoS classifiers*

Global QoS classifiers	Type of QoS marking used to prioritize outbound traffic	
	802.1p Priority[1] only	**DSCP policy[2]: DSCP codepoint with 802.1p priority**
UDP/TCP	Supported	Supported
IP Device	Supported	Supported
IP Precedence	Supported[3]	Not Supported
IP DiffServ	Supported	Supported
L3 Protocol	Supported	Not Supported
VLAN ID	Supported	Supported
Source Port	Supported	Supported

[1] When you configure only the 802.1p priority to mark packets that match a global QoS classifier, the selected traffic is prioritized and sent to the corresponding outbound port queue on the switch. VLAN-tagged ports are necessary to carry the 802.1p priority in a packet header to downstream devices.

[2] When you configure a DSCP policy to mark packets that match a global QoS classifier, the selected traffic is also prioritized according to the associated 802.1p priority and sent to the corresponding outbound port queue on the switch. VLAN-tagged ports carry the 802.1p priority in a packet header to downstream devices. In addition, you can configure downstream devices to read the DSCP value in IP packets and implement the service policy implied by the codepoint.

[3] When using a global QoS IP Precedence classifier, the 802.1p priority is automatically assigned to matching packets based on the IP precedence bit set in the packet header.

3. If you want 802.1p priority settings to be included in outbound packets, ensure that tagged VLANs are configured on the appropriate downstream links.
4. Determine the actual QoS configuration changes you will need to make on each QoS-capable device in your network in order to implement the desired policy. Also, if you want downstream devices to read and use DSCPs in IP packets from the switch, configure them to do so by enabling ToS Differentiated Service mode and making sure that the same DSCP policies are configured.

Using classifiers to configure QoS for outbound traffic

NOTE: In addition to the information in this section on the various QoS classifiers, see **QoS operating notes and restrictions** on page 213.

Viewing the QoS configuration

NOTE: UDP port priority and TCP port priority are not supported for IPv6 traffic on switches J9779A, J9780A, J9782A, and J9783A.

Examples of the `show qos` output are included with the example for each priority type.

Syntax:

```
show qos <priority-classifier>
```

`device-priority`: Displays the device priority table/configuration (priority based on the IP address).

`dscp-map`: Displays mappings between DSCP policy and 802.1p priority.

`port-priority`: Displays the current source-port priority configuration.

`protocol-priority`: Displays the protocol priority configuration.

`queue-config`: Displays the outbound port queue configuration information.

`resources`: Displays the resources used by the Policy Enforcement Engine.

NOTE: On switches J9779A, J9780A, J9782A, and J9783A, you may run out of QoS resources more quickly.

`tcp-udp-port-priority`: Displays the TCP/UDP port priorities.

`traffic-template`: Displays the configured traffic templates and their status.

`type-of-service`: Displays the current type-of-service priority configuration. The display output differs according to the ToS option used:

- `IP Precedence`
- `Diffserve`

`vlan-priority`: Displays the current VLAN priority configuration.

No override

By default, the `show` command outputs automatically list `No-override` for priority options that have not been configured. This means that if you do not configure a priority for a specific option, QoS does not prioritize packets to which that option applies, resulting in the `No-override` state.

- IP packets received through a VLAN-tagged port are managed using the 802.1p priority they carry in the 802.1Q field in their headers.
- VLAN-tagged packets received through an untagged port are handled by the switch with “normal” priority.

the `show qos vlan-priority` output shows the global QoS configurations on the switch that are configured with the VLAN ID classifier. Note that non-default 802.1p priorities have been configured for VLAN IDs 22 and 33; packets received on VLAN 1 are managed with the default settings, as described in the two bulleted items above.

Figure 26: *Output for the `show qos vlan-priority` command (example)*

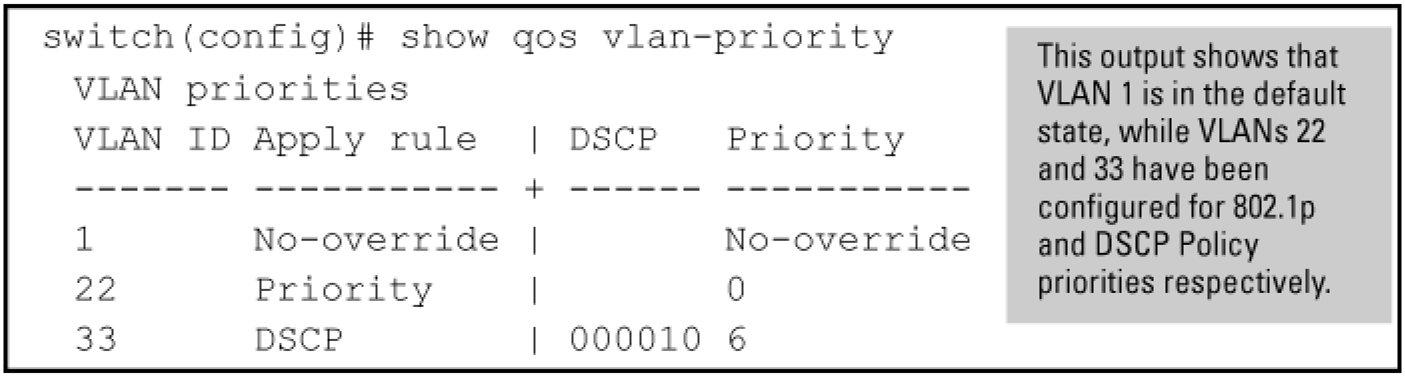

```
switch(config)# show qos vlan-priority
 VLAN priorities
 VLAN ID Apply rule  | DSCP   Priority
 ------- ----------- + ------ -----------
 1       No-override |        No-override
 22      Priority    |        0
 33      DSCP        | 000010 6
```

This output shows that VLAN 1 is in the default state, while VLANs 22 and 33 have been configured for 802.1p and DSCP Policy priorities respectively.

Global TCP/UDP classifier

Global QoS classifier precedence: 1

When you use TCP or UDP and a layer 4 Application port number as a global QoS classifier, traffic carrying the specified TCP/UDP port numbers is marked with a specified priority level, without regard for any other QoS classifiers in the switch.

You can configure up to 50 TCP/UDP application port numbers as QoS classifiers.

Options for assigning priority

The packet-marking options for global TCP/UDP port-number classifiers include:

- 802.1p priority
- DSCP policy: Assigning a new DSCP and 802.1p priority

For a given TCP or UDP port number, you can use only one of the above options at a time. However, for different port numbers, you can use different options.

TCP/UDP port number ranges

There are three ranges:

- Well-Known Ports: 0 – 1023
- Registered Ports: 1024 – 49151
- Dynamic and Private Ports: 49152 – 65535

For more information, including a listing of UDP/TCP port numbers, go to the **Internet Assigned Numbers Authority** (IANA) website at:

http://www.iana.org

Then click:

Protocol Number Assignment Services

P (under **Directory of General Assigned Numbers**)

Port Numbers

Assigning an 802.1p priority for a global TCP/UDP classifier

To mark matching TCP or UDP packets with an 802.1p priority, enter the following command:

Syntax:

```
qos < udp-port | tcp-port > [ ipv4 | ipv6 | ip-all ] < port-number | range start end > priority < 0-7>
```

Marks an 802.1p priority in outbound packets with the specified TCP or UDP application-port number, where:

`ipv4`: Marks only IPv4 packets (default).

`ipv6`: Marks only IPv6 packets.

`ip-all`: Marks all IP traffic (both IPv4 and IPv6 packets).

`port-number`: TCP/UDP port number from 1 to 65535.

`range <start end>` : Marks a range of TCP/UDP ports. If you specify a range, the minimum port number must precede the maximum port number in the range.

NOTE: Port range is not supported on switches J9779A, J9780A, J9782A, and J9783A.

`priority <0-7>`: Marks the specified 802.1p priority in matching TCP or UDP packets.

NOTE: UDP port priority and TCP port priority are not supported for IPv6 traffic on switches J9779A, J9780A, J9782A, and J9783A.

The 802.1p priority determines the packet's queue in the outbound port on the switch. If the packet leaves the switch on a tagged VLAN port, it carries the 802.1p priority with it to the next downstream device.

Default: Disabled — No 802.1p priority is assigned.

The `no` form of the command deletes the specified UDP or TCP port number or range of port numbers as a QoS classifier.

NOTE: If you have specified a range of port numbers, you must specify the entire range in the `no` command; you cannot remove part of a range.

Syntax:

```
show qos tcp-udp-port-priority
```

Displays a listing of all TCP and UDP QoS classifiers currently in the running-config file.

Operating notes on using TCP/UDP port ranges

- Only six concurrent policies are possible when using unique ranges. The number of policies allowed is less if ACLs are also using port ranges.
- No ranges allowed that include any port numbers configured as part of another QoS application port number policy.
- An error message is generated if there are not enough hardware resources available when configuring a policy.
- The entire range of configured port numbers must be specified when using the `no` form of the command, for example:

```
switch(config)# qos udp-port range 1300 1399 dscp 001110
switch(config)# no qos range 1300 1399
```

The following example displays the following configuration for TCP and UDP port prioritization:

Configuration for TCP and UDP port prioritization

TCP/UDP port	802.1p priority for TCP	802.1p priority for UDP
TCP Port 23 (Telnet)	7	7
UDP Port 23 (Telnet)	7	7

Table Continued

TCP/UDP port	802.1p priority for TCP	802.1p priority for UDP
TCP Port 80 (World Wide Web HTTP)	2	2
UDP Port 80 (World Wide Web HTTP)	1	1

Figure 27: *Configuring 802.1p priority assignments on TCP/UDP ports*

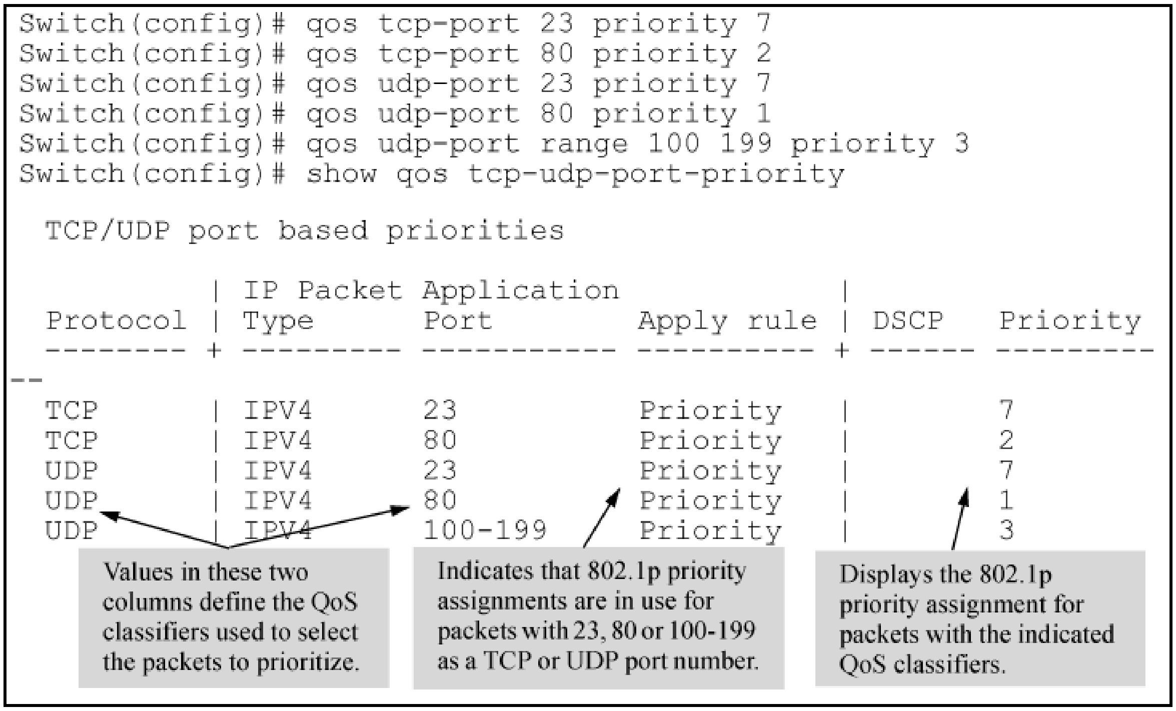

```
Switch(config)# qos tcp-port 23 priority 7
Switch(config)# qos tcp-port 80 priority 2
Switch(config)# qos udp-port 23 priority 7
Switch(config)# qos udp-port 80 priority 1
Switch(config)# qos udp-port range 100 199 priority 3
Switch(config)# show qos tcp-udp-port-priority

  TCP/UDP port based priorities

           | IP Packet Application            |
  Protocol | Type      Port        Apply rule | DSCP   Priority
  -------- + --------- ----------- ---------- + ------ ---------
--
  TCP      | IPV4      23          Priority   |        7
  TCP      | IPV4      80          Priority   |        2
  UDP      | IPV4      23          Priority   |        7
  UDP      | IPV4      80          Priority   |        1
  UDP      | IPV4      100-199     Priority   |        3
```

Values in these two columns define the QoS classifiers used to select the packets to prioritize.

Indicates that 802.1p priority assignments are in use for packets with 23, 80 or 100-199 as a TCP or UDP port number.

Displays the 802.1p priority assignment for packets with the indicated QoS classifiers.

Assigning a DSCP policy for a global TCP/UDP classifier

This global QoS packet-marking option assigns a previously configured or default DSCP policy (codepoint and 802.1p priority) to TCP or UDP packets having the specified port number or range of port numbers. When assigning a DSCP policy, the switch performs the following actions:

Procedure

1. Selects an incoming IP packet if the TCP or UDP port number it carries matches the port number specified in the TCP or UDP classifier (as shown in the figure in **Operating notes on using TCP/UDP port ranges**).
2. Overwrites (re-marks) the packet's DSCP with the new DSCP configured for matching packets.
3. Assigns the 802.1p priority associated with the new DSCP (see **Differentiated Services Codepoint (DSCP) mapping** on page 200).
4. Forwards the packet through the appropriate outbound port queue.

Creating a DSCP policy based on TCP/UDP port number classifiers

The following procedure creates a DSCP policy for IP packets carrying the selected TCP or UDP port-number classifier.

Procedure

1. Identify the TCP or UDP port-number classifier you want to use for assigning a DSCP policy.
2. Determine the DSCP policy for packets carrying the selected TCP or UDP port number or range of port numbers.
    a. Determine the DSCP you want to assign to the selected packets. (This codepoint will be used to overwrite (re-mark) the DSCP carried in packets received from upstream devices.)
    b. Determine the 802.1p priority you want to assign to the DSCP.
3. If necessary, use the `qos dscp-map` *`<codepoint>`* `priority <0-7>` command to configure the DSCP policy (codepoint and associated 802.1p priority) that you want to use to mark matching packets.

NOTE: Prerequisite: A DSCP codepoint must have a preconfigured 802.1p priority (0 - 7) before you can use the codepoint to mark matching packets. If a codepoint you want to use shows `No-override` in the Priority column of the DSCP Policy table (using the `show qos dscp-map` command), you must first configure a priority for the codepoint before proceeding (using the `qos dscp-map priority` command).

```
qos dscp-map <codepoint> priority <0-7>>
```

(Optional) This command is required only if an 802.1p priority is **not** already assigned to the specified *`<codepoint>`* in the DSCP Policy table. Valid values for a DSCP codepoint are as follows:

- A binary value for the six-bit codepoint from `000000` to `111111`.
- A decimal value from `0` (low priority) to `63` (high priority) that corresponds to a binary DSCP bit set
- An ASCII standard (hexadecimal) name for a binary DSCP bit set:

`af11` (001010)	`af42` (100100)
`af12` (001100)	`af43` (100110)
`af13` (001110)	`ef` (101110)
`af21` (010010)	`cs1` (001000) = precedence 1
`af22` (010100)	`cs2` (010000) = precedence 2
`af23` (010110)	`cs3` (011000) = precedence 3
`af31` (011010)	`cs4` (100000) = precedence 4
`af32` (011100)	`cs5` (101000) = precedence 5
`af33` (011110)	`cs6` (110000) = precedence 6

Table Continued

`af41` (100010)	`cs7` (111000) = precedence 7
`default` (000000)	

Enter **?** to display the list of valid codepoint entries.

When the switch applies the specified DSCP policy to a packet, the priority determines the packet's queue in the outbound port to which it is sent. If the packet leaves the switch on a tagged port, it carries the 802.1p priority with it to the next downstream device. For IP packets, the DSCP will be replaced by the codepoint specified in this command.

(Default: `No-override` for most codepoints.)

4. Configure the switch to assign the DSCP policy to packets with the specified TCP or UDP port number or range of port numbers.

```
no qos {udp-port | tcp-port} [ipv4 | ipv6 | ipv-all] {<port-number> | range
<port start><port end>} {dscp < codepoint> | priority <priority>}
```

```
no qos {udp-port | tcp-port} [ipv4 | ipv6 | ipv-all] {<port-number> | range
<port start><port end>} {dscp < codepoint> | priority <priority>}
```

Assigns a DSCP policy to outbound packets having the specified TCP or UDP application-port number or port range, and overwrites the DSCP in these packets with the assigned *`<codepoint>`* value, where:

- `port-number`: specifies a TCP/UDP port-number from 1 to 65535.
- `range` *`<start end>`*: specifies a range of TCP/UDP ports. If you specify a range, the minimum port number must precede the maximum port number in the range.
- `dscp` *`<codepoint>`*: overwrites the DSCP codepoint in the IPv4 ToS byte or IPv6 Traffic Class byte of matching packets with the specified value.Valid values for the DSCP codepoint are as follows:
  - A binary value for the 6-bit codepoint from `000000` to `111111`.
  - A decimal value from `0` (low priority) to `63` (high priority) that corresponds to a binary DSCP bit set
  - An ASCII standard name for a binary DSCP bit set

    Enter **?** to display the list of valid codepoint entries.

    The DSCP value you enter must be currently associated with an 802.1p priority in the DSCP Policy table. The 802.1p priority and determines the packet's queue in the outbound port to which it is sent. If the packet leaves the switch on a tagged port, it carries the 802.1p priority with it to the next downstream device.

    The default DSCP codepoint is `No-override`. The DSCP codepoint is not overwritten in matching packets.

The `no` form of the command deletes the specified UDP or TCP port number or range of port numbers as a QoS classifier. If you configured a range of port numbers as the QoS classifier, you must enter the entire range in the `no` command; you cannot remove part of a range.

Syntax

```
show qos tcp-udp-port-priority
```

Displays a listing of all TCP and UDP QoS classifiers currently in the running-config file.

Example:

This example shows how to assign the following DSCP policies to packets that match the specified TCP and UDP port applications:

Port Applications	DSCP Policies	
	DSCP	Priority
23-UDP	000111	7
80-TCP	000101	5
914-TCP	000010	1
1001-UDP	000010	1

1. Determine if the DSCP codepoints that you want to use to mark matching packets already have an 802.1p priority assigned, which could indicate use by existing applications (`show qos dscp-map` command).

   A DSCP codepoint must also have a priority configured before you can use it to mark matching packets.

```
switch(config)# show qos dscp-map

  DSCP -> 802.p priority mappings

  NOTE: ‘qos type-of-service diff-services’ must be configured
         before DSCP is honored on inbound traffic.

  DSCP CodePoint DSCP Value 802.1p tag   DSCP Policy name
  -------------- ---------- -----------  ----------------------
  000000         0          0            cs0
  000001         1          No-override
  000010         2          No-override
  000011         3          No-override
  000100         4          No-override
  000101         5          No-override
  000110         6          No-override
  000111         7          No-override
  001000         8          1            cs1
  001001         9          No-override
```

2. Configure the DSCP policies for the codepoints you want to use.

```
switch(config)# qos dscp-map af11 priority 3
switch(config)# qos dscp-map 13 priority 3
switch(config)# qos dscp-map af13 priority 3
switch(config)# write memory

switch(config)# show config
switch configuration:

; J9146 Configuration Editor; Created on release XX.15.XX

hostname “Switch”
time daylight-time-rule None
```

```
qos dscp-map af11 priority 3
qos dscp-map 13 priority 3
qos dscp-map af13 priority 3
...
```

3. Assign the DSCP policies to the selected TCP/UDP port applications and display the result.

```
switch(config)# qos udp-port 23 dscp 000111
switch(config)# qos tcp-port 80 dscp 000101
switch(config)# qos tcp-port 914 dscp 000010
switch(config)# qos udp-port range 1001 2000 dscp 000010

  TCP/UDP port based priorities

           | IP Packet Application            |
  Protocol | Type      Port         Apply rule | DSCP   Priority
  -------- + --------- ----------- ---------- + ------ -----------
  UDP      | IPV4      23          DSCP       | 8      7
  TCP      | IPV4      80          DSCP       | 6      5
  TCP      | IPV4      914         DSCP       | 3      1
  UDP      | IPV4      1001-2000   DSCP       | 3      1
```

The switch applies the DSCP policies in the above output to IP packets with the specified TCP/UDP port applications that are received in the switch. The switch manages the packets as follows:

- Overwrites the original DSCPs in the selected packets with the new DSCPs specified in the above policies.
- Assigns the 802.1p priorities in the above policies to the selected packets.

Global IP-device classifier

Global QoS classifier precedence: 2

The global IP-device classifier enables you to configure up to 250 IP addresses to select IP packets according to source **or** destination address.

NOTE: QoS IP-device restriction: The switch does not allow a QoS IP-device priority for the Management VLAN IP address (if configured). If no Management VLAN is configured, then the switch does not allow configuring a QoS IP-device priority for the default VLAN IP address.

Options for assigning priority

The packet-marking options for global IP-device classifiers include:

- 802.1p priority
- DSCP policy: Assigning a new DSCP and 802.1p priority

For a given IP address or subnet mask, you can assign only one of the above options at a time. However, for different IP addresses, you can use different options.

QoS IP Type-of-Service (ToS) policy and priority

Global QoS classifier precedence: 3

You can assign a maximum of 64 ToS rules. This feature applies only to IPv4 traffic and performs either of the following:

- **ToS IP-precedence mode:** All IP packets generated by upstream devices and applications include precedence bits in the ToS byte. Using this mode, the switch uses these bits to compute and assign the corresponding 802.1p priority.
- **ToS Differentiated Services (Diffserv) mode:** This mode requires knowledge of the codepoints set in IP packets by the upstream devices and applications. It uses the ToS codepoint in IP packets coming from upstream devices and applications to assign 802.1p priorities to the packets. You can use this option to do both of the following:
    - **Assign a new prioritization policy:** A "policy" includes both a codepoint and a corresponding 802.1p priority. This option selects an incoming IPv4 packet on the basis of its codepoint and assigns a new codepoint and corresponding 802.1p priority. (Use the `qos dscp-map` command to specify a priority for any codepoint; see **Differentiated Services Codepoint (DSCP) mapping** on page 200.)
    - **Assign an 802.1p priority:** This option reads the DSCP of an incoming IPv4 packet and, without changing this codepoint, assigns the 802.1p priority to the packet, as configured in the DSCP Policy Table (**Differentiated Services Codepoint (DSCP) mapping** on page 200). This means that a priority value of 0 – 7 must be configured for a DSCP before the switch will attempt to perform a QoS match on the packet's DSCP bits.

  Before configuring the ToS Diffserv mode, you must use the `qos dscp-map` command to configure the desired 802.1p priorities for the codepoints you want to use for either option. See **Differentiated Services Codepoint (DSCP) mapping** on page 200 for more information.

Unless IP-Precedence mode and Diffserv mode are both disabled (the default setting), enabling one automatically disables the other.

NOTE: "Mixing" ToS DSCP policies and 802.1p priorities is not recommended.

Assigning an 802.1p priority to IPv4 packets on the basis of the ToS precedence bits

If a device or application upstream of the switch sets the precedence bits in the ToS byte of IPv4 packets, you can use this feature to apply that setting for prioritizing packets for outbound port queues. If the outbound packets are in a tagged VLAN, this priority is carried as an 802.1p value to the adjacent downstream devices.

Syntax:

```
qos type-of-service ip-precedence
```

Causes the switch to automatically assign an 802.1p priority to all IPv4 packets by computing each packet's 802.1p priority from the precedence bits the packet carries. This priority determines the packet's queue in the outbound port to which it is sent. If the packet leaves the switch on a tagged port, it carries the 802.1p priority with it to the next downstream device.

(ToS IP Precedence Default: Disabled)

```
no qos type-of-service
```

Disables all ToS classifier operation, including prioritization using the precedence bits.

```
show qos type-of-service
```

When the IP-precedence mode is enabled (or if neither ToS option is configured), this command displays the ToS configuration status. If the Diff-serv mode is enabled, codepoint data is displayed.

Using the IP-precedence classifier, prioritization of outbound packets relies on the IP-Precedence bit setting that IP packets carry with them from upstream devices and applications. To configure and verify this option:

Figure 28: *Enabling ToS IP-precedence prioritization*

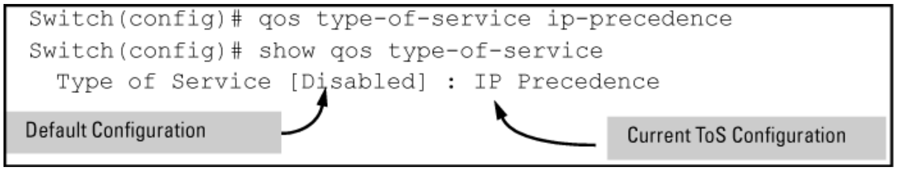

```
Switch(config)# qos type-of-service ip-precedence
Switch(config)# show qos type-of-service
  Type of Service [Disabled] : IP Precedence
```

To replace this option with the ToS diff-services option, configure diff-services as described below, which automatically disables IP-Precedence. To disable IP-Precedence without enabling the diff-services option, use this command: `no qos type-of-service`

Assigning an 802.1p priority to IPv4 packets on the basis of incoming DSCP

One of the best uses for this option is on an interior switch where you want to honor (continue) a policy set on an edge switch. That is, it enables you to select incoming packets having a specific DSCP and forward these packets with the desired 802.1p priority. For example, if an edge switch “A” marks all packets received on port 5 with a particular DSCP, you can configure a downstream (interior) switch “B” to handle such packets with the desired priority (regardless of whether 802.1Q tagged VLANs are in use).

Figure 29: *Interior switch B honors the policy established in edge switch A*

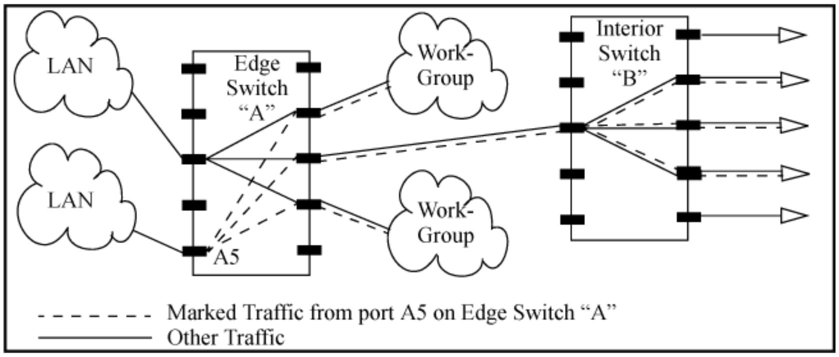

To do so, assign the desired 802.1p priority to the same codepoint that the upstream or edge switch assigns to the selected packets. When the downstream switch receives an IPv4 packet carrying one of these codepoints, it assigns the configured priority to the packet and sends it out the appropriate priority queue. (The packet retains the codepoint it received from the upstream or edge switch). You can use this option concurrently with the diffserv DSCP Policy option (described later in this section), as long as the DSCPs specified in the two options do not match.

NOTE: Regarding DSCP use: Different applications may use the same DSCP in their IP packets. Also, the same application may use multiple DSCPs if the application originates on different clients, servers, or other devices. Using an edge switch enables you to select the desired packets and mark them with predictable DSCPs that can be used by downstream switches to honor policies set in the edge switch.

When enabled, the switch applies direct 802.1p prioritization to all packets having codepoints that meet these prerequisites:

- The codepoint is configured with an 802.1p priority in the DSCP table. (Codepoints configured with `No-override` are not used.)
- The codepoint is not configured for a new DSCP policy assignment.

Thus, the switch does not allow the same incoming codepoint (DSCP) to be used simultaneously for directly assigning an 802.1p priority and also assigning a DSCP policy. For a given incoming codepoint, if you configure one option and then the other, the second overwrites the first.

To use this option:

Procedure

1. Identify a DSCP used to set a policy in packets received from an upstream or edge switch.
2. Determine the 802.1p priority (0–7) you want to apply to packets carrying the identified DSCP. (You can either maintain the priority assigned in the upstream or edge switch, or assign a new priority.)
3. Use `qos dscp-map <codepoint> priority <0-7>` to assign the 802.1p priority you want to the specified DSCP.
4. Enable `diff-services` if not already enabled.

Syntax:

```
qos type-of-service diff-services <codepoint>
```

Causes the switch to read the *<codepoint>* (DSCP) of an incoming IPv4 packet and, when a match occurs, assign a corresponding 802.1p priority, as configured in the switch's DSCP table (see **Differentiated Services Codepoint (DSCP) mapping**).

```
no qos type-of-service
```

Disables all ToS classifier operation.

```
no qos dscp-map <codepoint>
```

Disables direct 802.1p priority assignment to packets carrying the *<codepoint>* by reconfiguring the codepoint priority assignment in the DSCP table to `No-override`. If this codepoint is in use as a DSCP policy for another diffserv codepoint, you must disable or redirect the other diffserv codepoint's DSCP policy before you can disable or change the codepoint. For example, in **Figure 31: ToS configuration that enables both 802.1p priority and DSCP policy assignment** on page 187 you cannot change the priority for the 000000 codepoint until you redirect the DSCP policy for 000001 away from using 000000 as a policy. (See **Note on changing a priority setting** on page 202 and **Differentiated Services Codepoint (DSCP) mapping** on page 200.)

```
show qos type-of-service
```

Displays current Type-of-Service configuration. In diffserv mode it also shows the current direct 802.1p assignments and the current DSCP assignments covered later in this section.

For example, an edge switch “A” in an untagged VLAN assigns a DSCP of 000110 on IP packets it receives on port 6, and handles the packets with high priority (7). When these packets reach interior switch “B” you want the switch to handle them with the same high priority. To enable this operation, you would configure an 802.1p priority of 7 for packets received with a DSCP of 000110. ToS `diff-services` must be enabled as shown in the following images.

Figure 30: *Viewing the codepoints available for 802.1p priority assignments*

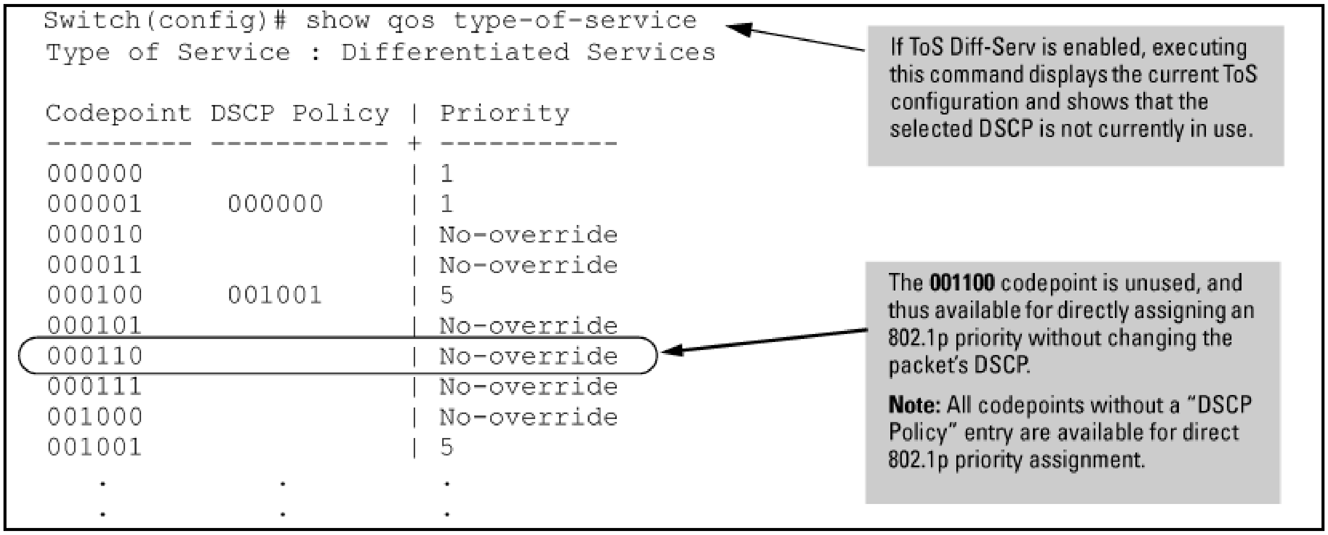

```
Switch(config)# show qos type-of-service
Type of Service : Differentiated Services

Codepoint DSCP Policy | Priority
--------- ----------- + -----------
000000                | 1
000001     000000     | 1
000010                | No-override
000011                | No-override
000100     001001     | 5
000101                | No-override
000110                | No-override
000111                | No-override
001000                | No-override
001001                | 5
   .           .        .
   .           .        .
```

If ToS Diff-Serv is enabled, executing this command displays the current ToS configuration and shows that the selected DSCP is not currently in use.

The **001100** codepoint is unused, and thus available for directly assigning an 802.1p priority without changing the packet's DSCP.

Note: All codepoints without a “DSCP Policy” entry are available for direct 802.1p priority assignment.

Figure 31: *ToS configuration that enables both 802.1p priority and DSCP policy assignment*

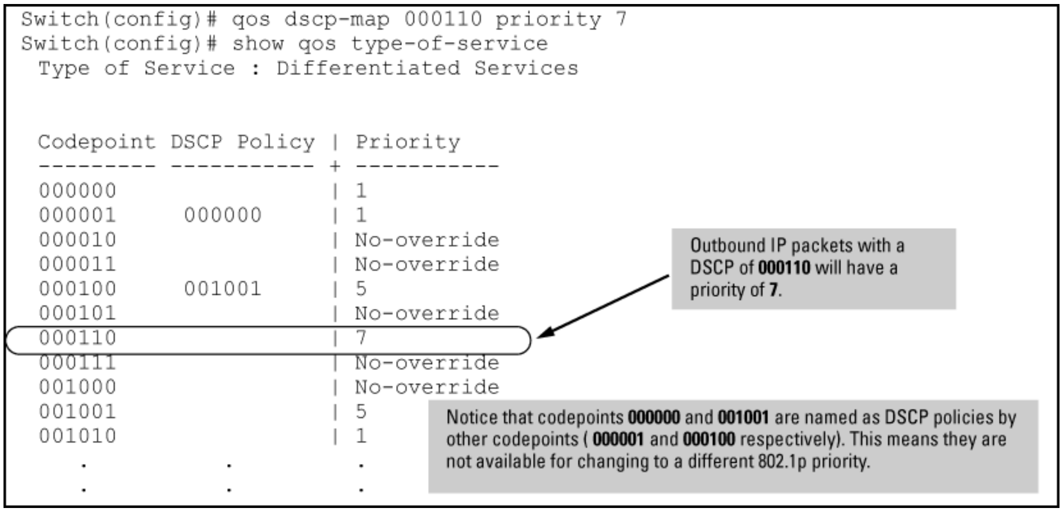

```
Switch(config)# qos dscp-map 000110 priority 7
Switch(config)# show qos type-of-service
  Type of Service : Differentiated Services


  Codepoint DSCP Policy | Priority
  --------- ----------- + -----------
  000000                | 1
  000001     000000     | 1
  000010                | No-override
  000011                | No-override
  000100     001001     | 5
  000101                | No-override
  000110                | 7
  000111                | No-override
  001000                | No-override
  001001                | 5
  001010                | 1
     .           .        .
     .           .        .
```

Outbound IP packets with a DSCP of **000110** will have a priority of **7**.

Notice that codepoints **000000** and **001001** are named as DSCP policies by other codepoints (**000001** and **000100** respectively). This means they are not available for changing to a different 802.1p priority.

Assigning a DSCP policy on the basis of the DSCP in IPv4 packets received from upstream devices

The preceding section describes how to forward a policy set by an edge (or upstream) switch. This option changes a DSCP policy in an IPv4 packet by changing its IP ToS codepoint and applying the priority associated with the new codepoint. (A DSCP policy consists of a differentiated services codepoint and an associated 802.1p priority.) You can use this option concurrently with the diffserv 802.1p priority option (above), as long as the DSCPs specified in the two options do not match.

To use this option to configure a change in policy:

Procedure

1. Identify the DSCP used to set a policy in packets received from an upstream or edge switch.
2. Create a new policy by using the `qos dscp-map <code-point> priority <0-7>` command to configure an 802.1p priority for the codepoint you will use to overwrite the DSCP that the packet carries from upstream.
3. Use the `qos type-of-service diff-services < mapped to DSCP > dscp < mapped from DSCP >` command to change the policy on packets coming from the edge or upstream switch with the specified incoming DSCP.

   Figure 29: Interior switch B honors the policy established in edge switch A on page 185 illustrates this scenario

Syntax:

```
qos type-of-service diff-services
```

Enables ToS Diff-serve QoS so that Diff-serve policy configurations can be applied to incoming packets that have matching codepoints.

Syntax:

```
qos type-of-service diff-services <current-codepoint> dscp <new-codepoint>
```

Configures the switch to select an incoming IP packet carrying the *<current-codepoint>* and then use the *<new-codepoint>* to assign a new, previously configured DSCP policy to the packet. The policy overwrites the *<current-codepoint>* with the *<new-codepoint>* and assigns the 802.1p priority specified by the policy.

Syntax:

```
no qos type-of-service
```

Disables all ToS classifier operation. Current ToS DSCP policies and priorities remain in the configuration and will become available if you re-enable ToS Diff-services.

Syntax:

```
no qos type-of-service [diff-services <codepoint>]
```

Deletes the DSCP policy assigned to the *<codepoint>* and returns the *<codepoint>* to the 802.1p priority setting it had before the DSCP policy was assigned, which is either a value from 0 - 7 or `No-override`.

Syntax:

```
show qos type-of-service
```

Displays a listing of codepoints with any corresponding DSCP policy reassignments for outbound packets. Also displays the 802.1p priority for each codepoint that does not have a DSCP remarking policy assigned to it.

Example

For example, suppose that you want to configure the following two DSCP policies for packets received with the indicated DSCPs.

Received DSCP	Policy DSCP	802.1p Priority	Policy Name (Optional)
001100	17	6	Level 6
001101	16	4	Level 4

1. Determine whether the DSCPs already have priority assignments, which could indicate use by existing applications. This is not a problem as long as the configured priorities are acceptable for all applications using the same DSCP. (See **Note on changing a priority setting** on page 202. Also, a DSCP must have a priority configured before you can assign any QoS classifiers to use it. See **Differentiated Services Codepoint (DSCP) mapping** on page 200.)
2. After configuring the DSCP policies for the codepoints you want to use, assign the policies to the codepoints in the selected packet type.

   An example of policy assignment to outbound packets on the basis of the DSCP in the packets received from upstream devices is shown below. The specified DSCP policies overwrite the original DSCPs on the selected packets, and use the 802.1p priorities previously configured.

```
switch(config)# qos type-of-service diff-services 001100 dscp 17
switch(config)# qos type-of-service diff-services 001101 dscp 16
switch(config)# show qos type-of-service
  Type of Service : Differentiated Services

  Codepoint DSCP Policy | Priority
  --------- ----------- + -----------
  000000                | No-override
  000001                | No-override
  000010                | No-override
  000011                | No-override
  000100                | No-override
  000101                | No-override
  000110                | No-override
  000111                | No-override
  001000    001011      | 7
  001001                | No-override
  001010                | 1
  001011                | 7
  001100    010001      | 6
  001101    010000      | 4
```

Details of QoS IP ToS

IP packets include a Type of Service (ToS) byte. The ToS byte includes:

Precedence Bits: This element is a subset of the DSCP and is composed of the upper 3 bits of the ToS byte. When configured to do so, the switch uses the precedence bits to determine a priority for handling the associated packet. (The switch does not change the setting of the precedence bits.) Using the ToS Precedence bits to prioritize IPv4 packets relies on priorities set in upstream devices and applications.

The following figure, shows an example of the ToS byte in the header for an IPv4 packet, and illustrates the diffserv bits and precedence bits in the ToS byte. (Note that the Precedence bits are a subset of the Differentiated Services bits.)

Figure 32: *The ToS codepoint and precedence bits*

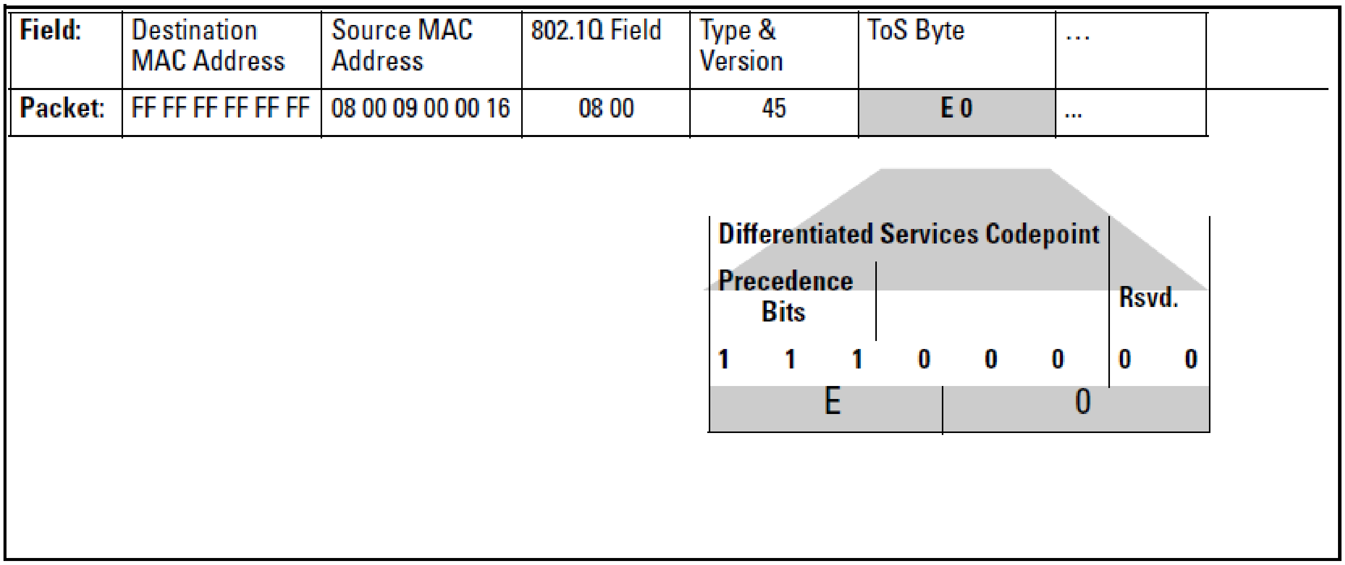

Field:	Destination MAC Address	Source MAC Address	802.1Q Field	Type & Version	ToS Byte	…
Packet:	FF FF FF FF FF FF	08 00 09 00 00 16	08 00	45	**E 0**	...

Table 22: *How the switch uses the ToS configuration*

Outbound port	ToS option:	
	802.1p (value = 0 - 7)	**Differentiated services**
IP packet sent out an untagged port in a VLAN	Depending on the value of the IP Precedence bits in the packet's ToS field, the packet will go to one of eight outbound port queues in the switch. See the table in **Overview of QoS settings**.	For a given packet carrying a ToS codepoint that the switch has been configured to detect: • Change the codepoint according to the configured policy and assign the 802.1p priority specified for the new codepoint in the DSCP Policy Table (**Differentiated Services Codepoint (DSCP) mapping**). • Do not change the codepoint, but assign the 802.1p priority specified for the existing codepoint in the DSCP Policy Table (**Differentiated Services Codepoint (DSCP) mapping**). Depending on the 802.1p priority used, the packet will leave the switch through a queue as defined in the table in **Overview of QoS settings** on page 170.If `No-override` (the default) has been configured for a specified codepoint, then the packet is not prioritized by ToS and, by default, is sent to the "normal priority" queue.
IP packet sent out an untagged port in a VLAN	Same as above, plus the IP Precedence value (0 - 7) will be used to set a corresponding 802.1p priority in the VLAN tag carried by the packet to the next downstream device. See the table below.	Same as above, plus the Priority value (0 - 7) will be used to set a corresponding 802.1p priority in the VLAN tag carried by the packet to the next downstream device. Where `No-override` is the assigned priority, the VLAN tag carries a "0" (normal priority) 802.1p setting if not prioritized by other QoS classifiers.

Table 23: *ToS IP-precedence bit mappings to 802.1p priorities*

ToS byte IP precedence bits	Corresponding 802.1p priority	Service priority level
000	1	Lowest
001	2	Low
002	0	Normal
003	3	

Table Continued

ToS byte IP precedence bits	Corresponding 802.1p priority	Service priority level
004	4	
005	5	
006	6	
007	7	Highest

Global Layer-3 protocol classifier

Global QoS classifier precedence: 4

When a global Layer-3 Protocol classifier is configured as the highest-precedence classifier and the switch receives traffic carrying the specified protocol, matching packets are assigned the priority configured for the classifier.

Assigning a priority for a global Layer-3 protocol classifier

This global QoS packet-marking option assigns an 802.1p priority to outbound packets having the specified Layer-3 protocol.

Syntax:

```
qos protocol < ip | ipx | arp | appletalk | sna | netbeui> priority < 0 - 7 >
```

Configures an 802.1p priority for outbound packets having the specified protocol. This priority determines the packet's queue in the outbound port to which it is sent. If the packet leaves the switch on a tagged port, it carries the 802.1p priority with it to the next downstream device. You can configure one QoS classifier for each protocol type.

(Default: `No-override`)

Syntax:

```
no qos protocol < ip | ipx | arp | appletalk | sna | netbeui >
```

Disables use of the specified protocol as a QoS classifier and resets the protocol priority to `No-override`.

Syntax:

```
show qos protocol-priority
```

Lists the QoS protocol classifiers with their priority settings.

Configuring global Layer-3 protocol classifiers

To configure the following global Layer-3 protocol classifiers:

1. Configure QoS protocol classifiers with IP at 0 (normal), ARP at 5 (medium), and AppleTalk at 7 (high) and display the QoS protocol configuration.
2. Disable the QoS IP protocol classifier, downgrade the ARP priority to 4, and again display the QoS protocol configuration.

The following example shows the necessary configuration commands.

Figure 33: *Adding, viewing, removing, and changing QoS protocol classifiers*

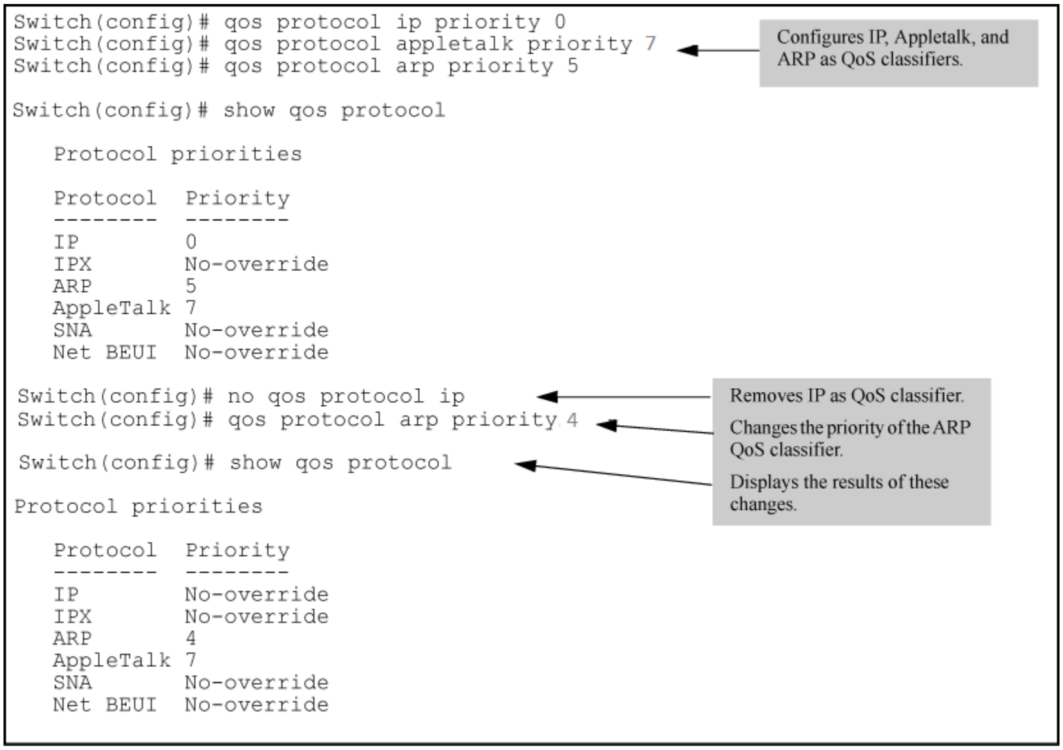

```
Switch(config)# qos protocol ip priority 0
Switch(config)# qos protocol appletalk priority 7
Switch(config)# qos protocol arp priority 5

Switch(config)# show qos protocol

   Protocol priorities

   Protocol  Priority
   --------  --------
   IP        0
   IPX       No-override
   ARP       5
   AppleTalk 7
   SNA       No-override
   Net BEUI  No-override

Switch(config)# no qos protocol ip
Switch(config)# qos protocol arp priority 4

Switch(config)# show qos protocol

Protocol priorities

   Protocol  Priority
   --------  --------
   IP        No-override
   IPX       No-override
   ARP       4
   AppleTalk 7
   SNA       No-override
   Net BEUI  No-override
```

Configures IP, Appletalk, and ARP as QoS classifiers.

Removes IP as QoS classifier.

Changes the priority of the ARP QoS classifier.

Displays the results of these changes.

QoS VLAN-ID (VID) priority

Global QoS classifier precedence: 5

Where a particular VLAN-ID classifier has the highest precedence in the switch for traffic in that VLAN, then traffic received in that VLAN is marked with the VID classifier's configured priority level. Different VLAN-ID classifiers can have differing priority levels.

Options for assigning priority

Priority control options for packets carrying a specified VLAN-ID include:

- 802.1p priority
- DSCP policy: Assigning a new DSCP and 802.1p priority

(For operation when other QoS classifiers apply to the same traffic, see **Classifiers for prioritizing outbound packets** on page 172.)

NOTE: QoS with VID priority applies to static VLANs only, and applying QoS to dynamic VLANs created by GVRP operation is not supported. A VLAN must exist while a subject of a QoS configuration, and eliminating a VLAN from the switch causes the switch to clear any QoS features configured for that VID.

Assigning a priority based on VLAN-ID

This option assigns a priority to all outbound packets having the specified VLAN-ID (VID). You can configure this option by either specifying the VID ahead of the `qos` command or moving to the VLAN context for the VLAN you want to configure for priority.

Syntax:

```
vlan <vid> qos priority <0-7>
```

Configures an 802.1p priority for outbound packets belonging to the specified VLAN. This priority determines the packet's queue in the outbound port to which it is sent. If the packet leaves the switch on a tagged port, it carries the 802.1p priority with it to the next downstream device. You can configure one QoS classifier for each VLAN-ID.

(Default: `No-override`)

Syntax:

```
no vlan <vid> qos
```

Removes the specified VLAN-ID as a QoS classifier and resets the priority for that VLAN to `No-override`.

Syntax:

```
show qos vlan-priority
```

Displays a listing of the QoS VLAN-ID classifiers currently in the running-config file, with their priority data.

Procedure

1. For example, suppose that you have the following VLANs configured on the switch and want to prioritize them as shown:

```
switch(config)# show vlan
 Status and Counters - VLAN Information

  Maximum VLANs to support : 8
  Primary VLAN : DEFAULT_VLAN

  802.1Q VLAN ID Name         Status
  -------------- ----------   ------------
  1              DEFAULT_VLAN static
  22             VLAN_22      static
```

2. You would then execute the following commands to prioritize the VLANs by VID:

```
switch(config)# vlan 1 qos dscp 9
switch(config)# vlan 22 qos dscp 8

switch(config)# show qos vlan-priority

  VLAN priorities

  VLAN ID Apply rule  | DSCP   Priority
  ------- ----------- + ------ -----------
  1       DSCP        | 001001 7
  22      DSCP        | 001000 6
```

3. If you then decided to remove VLAN_22 from QoS prioritization:

   In this instance, `No-override` indicates that VLAN 22 is not prioritized by QoS.

```
switch(config)# no vlan 22 qos
switch(config)# show qos vlan

  VLAN priorities
  VLAN ID Apply rule   | DSCP   Priority
  ------- ----------- + ------ -----------
  1       DSCP         | 001001 7
  22      No-override |        No-override
```

Assigning a DSCP policy based on VLAN-ID

This option assigns a previously configured DSCP policy (codepoint and 802.1p priority) to outbound IP packets having the specified VLAN-ID (VID). The switch performs the following:

Procedure

1. Selects an incoming IP packet on the basis of the VLAN-ID it carries.
2. Overwrites the packet's DSCP with the DSCP configured in the switch for such packets.
3. Assigns 802.1p priority configured in the switch for the new DSCP (see **Differentiated Services Codepoint (DSCP) mapping** on page 200).
4. Forwards the packet through the appropriate outbound port queue.

Steps for creating a policy based on VLAN-ID classifier:

1. Determine the VLAN-ID classifier to which you want to assign a DSCP policy.
2. Determine the DSCP policy for packets carrying the selected VLAN-ID:
   a. Determine the DSCP you want to assign to the selected packets. (This codepoint will be used to overwrite the DSCP carried in packets received through the source-port from upstream devices.)
   b. Determine the 802.1p priority you want to assign to the DSCP.
3. Configure the DSCP policy by using `qos dscp-map` to configure the priority for each codepoint (see **Differentiated Services Codepoint (DSCP) mapping** on page 200 for more information).
4. Configure the switch to assign the DSCP policy to packets with the specified VLAN-ID.

Syntax:

```
vlan <vid> qos dscp <codepoint>
```

Assigns a DSCP policy to packets carrying the specified VLAN-ID, and overwrites the DSCP in these packets with the assigned `<codepoint>` value. This policy includes an 802.1p priority and determines the packet's queue in the outbound port to which it is sent. If the packet leaves the switch on a tagged port, it carries the 802.1p priority with downstream device.

(Default: `No-override`)

Syntax:

```
no vlan <vid> qos
```

Removes QoS classifier for the specified VLAN.

Syntax:

```
show qos vlan-priority
```

Displays a listing of the QoS VLAN-ID classifiers currently in the running-config file.

For example, suppose that you wanted to assign this set of priorities:

VLAN-ID	DSCP	Priority
40	15	7
30	16	5
20	17	1
1	17	1

Assign the DSCP policies to the selected VIDs and display the result.

An example of the completed VID-DSCP priority configuration is shown below.

```
switch(config)# vlan 1 qos dscp 17
switch(config)# vlan 20 qos dscp 17
switch(config)# vlan 30 qos dscp 16
switch(config)# vlan 40 qos dscp 15

switch(config)# show qos vlan-priority

 VLAN priorities

  VLAN ID Apply rule  | DSCP   Priority
  ------- ----------- + ------ -----------
  1       DSCP        | 010001 1
  20      DSCP        | 010001 1
  30      DSCP        | 010000 5
  40      DSCP        | 001111 7
```

In the example above, the switch will now apply the DSCP policies to packets received on the switch with the specified VLAN-IDs. This means the switch will:

- Overwrite the original DSCPs in the selected packets with the new DSCPs specified in the above policies.
- Assign the 802.1p priorities in the above policies to the appropriate packets.

QoS source-port priority

Global QoS classifier precedence: 6

The QoS source-port option enables you to use a packet’s source-port on the switch as a QoS classifier. Where a particular source-port classifier has the highest precedence in the switch for traffic entering through that port, then traffic received from the port is marked with the source-port classifier’s configured priority level. Different source-port classifiers can have different priority levels.

Options for assigning priority on the switch

Priority control options for packets from a specified source-port include:

- 802.1p priority
- DSCP policy: Assigning a new DSCP and 802.1p priority

(For operation when other QoS classifiers apply to the same traffic, see **Classifiers for prioritizing outbound packets** on page 172.)

Options for assigning priority from a RADIUS server

You can use a RADIUS server to impose a QoS source-port priority during an 802.1X port-access authentication session. See the RADIUS chapter in the *ArubaOS-Switch Access Security Guide* for your switch.

Assigning a priority based on source-port

This option assigns a priority to all outbound packets having the specified source-port. You can configure this option by either specifying the source-port ahead of the `qos` command or moving to the port context for the port you want to configure for priority. (If you are configuring multiple source-ports with the same priority, you may find it easier to use the `interface <port-list>` command to go to the port context instead of individually configuring the priority for each port.)

Syntax:

```
interface <port-list> qos priority <0-7>
```

Configures an 802.1p priority for packets entering the switch through the specified (source) ports. This priority determines the packet queue in the outbound ports to which traffic is sent. If a packet leaves the switch on a tagged port, it carries the 802.1p priority with it to the next downstream device. You can configure one QoS classifier for each source-port or group of source-ports.

(Default: `No-override`)

Syntax:

```
no interface <port-list> qos
```

Disables use of the specified source-ports for QoS classifiers and resets the priority for the specified sourceports to `No-override`.

Syntax:

```
show qos port-priority
```

Lists the QoS port-priority classifiers with their priority data.

For example, suppose that you want to prioritize inbound traffic on the following source-ports:

Source-port	Priority
1–3	2
4	3

You would then execute the following commands to prioritize traffic received on the above ports.

Configuring and displaying source-port QoS priorities

```
switch(config)# interface e 1-3 qos priority 2
switch(config)# interface e 4 qos priority 3
switch(config)# show qos port-priority

  Port priorities
  Port Apply rule  | DSCP  Priority    Radius Override
  ---- ----------    ----  --------    ---------------
  1    Priority    |       2           No-override
  2    Priority    |       2           No-override
  3    Priority    |       2           No-override
  4    Priority    |       3           No-override
  5    No-override |       No-override No-override
  .    .                   .           .
  .    .                   .           .
```

If you then decided to remove port 1 from QoS prioritization:

Returning a QoS-prioritized VLAN to "No-override" status

In this instance, `No-override` indicates that port 1 is not prioritized by QoS.

```
switch(config)# no interface 1 qos
switch(config)# show qos port-priority

  Port priorities

  Port Apply rule  | DSCP  Priority    Radius Override
  ---- ----------    ----  --------    ---------------
  1    No-override |       No-override No-override
  2    Priority    |       2           No-override
  3    Priority    |       2           No-override
  4    Priority    |       3           No-override
  5    No-override |       No-override No-override
  .    .                   .           .
  .    .                   .           .
```

Assigning a DSCP policy based on the source-port

This option assigns a previously configured DSCP policy (codepoint and 802.1p priority) to outbound IP packets (received from the specified sourceports). That is, the switch:

Procedure

1. Selects an incoming IP packet on the basis of its source-port on the switch.
2. Overwrites the packet's DSCP with the DSCP configured in the switch for such packets.
3. Assigns 802.1p priority configured in the switch for the new DSCP (see **Differentiated Services Codepoint (DSCP) mapping** on page 200).
4. Forwards the packet through the appropriate outbound port queue.

Steps for creating a policy based on source-port classifier:

NOTE: You can select one DSCP per source-port. Also, configuring a new DSCP for a source-port automatically overwrites (replaces) any previous DSCP or 802.1p priority configuration for that port.)

1. Identify the source-port classifier to which you want to assign a DSCP policy.
2. Determine the DSCP policy for packets having the selected source-port:

a. Determine the DSCP you want to assign to the selected packets. (This codepoint will be used to overwrite the DSCP carried in packets received through the source-port from upstream devices.)

b. Determine the 802.1p priority you want to assign to the DSCP.

3. Configure the DSCP policy by using `qos dscp-map` to configure the priority for each codepoint (see **Differentiated Services Codepoint (DSCP) mapping** on page 200 for more information).
4. Configure the switch to assign the DSCP policy to packets from the specified source-port.

Syntax:

```
interface <port-list> qos dscp <codepoint>
```

Assigns a DSCP policy to packets from the specified sourceports, and overwrites the DSCP in these packets with the assigned *<codepoint>* value. This policy includes an 802.1p priority and determines the packet's queue in the outbound port to which it is sent. If the packet leaves the switch on a tagged port, it carries the 802.1p priority with it to the next downstream device.

(Default: `No-override`)

Syntax:

```
no interface <port-list> qos
```

Removes QoS classifier for the specified source-ports.

Syntax:

```
show qos port
```

Displays a listing of all source-port QoS classifiers currently in the running-config file.

For example, suppose that you wanted to assign this set of priorities that have been configured on the switch:

Source-port	DSCP	Priority
2	15	7
1, 3	16	5
4, 5	17	1

Assign the DSCP policies to the selected source-ports and display the result.

An example of the completed source-port DSCP-priority configuration is shown below

```
switch(config)# int 4,5
switch(eth-4,5)# qos dscp 17
switch(eth-4,5)# int 1,3
switch(eth-1,3)# qos dscp 16
switch(eth-1,3)# int 2
switch(eth-2)# qos dscp 15

switch(eth-2)# show qos port-priority

Port priorities
Port Apply rule  | DSCP   Priority    Radius Override
---- ----------- + ------ ----------- ---------------
1    DSCP        | 010000 5           No-override
2    DSCP        | 001111 7           No-override
3    DSCP        | 010000 5           No-override
4    DSCP        | 010001 1           No-override
5    DSCP        | 010001 1           No-override
6    No-override |        No-override No-override
7    No-override |        No-override No-override
.         .         .          .           .
.         .         .          .           .
```

RADIUS override field

During a client session authenticated by a RADIUS server, the server can impose a port priority that applies only to that client session. For more information, see the RADIUS chapter in the *ArubaOS-Switch Access Security Guide* for your switch.

Differentiated Services Codepoint (DSCP) mapping

A codepoint must have an 802.1p priority assignment (0 - 7) before you can configure a policy for prioritizing packets. If a codepoint you want to use shows `No-override` in the `Priority` column of the DSCP map (`show qos dscp-map`), then you must assign a 0 - 7 priority before proceeding (`qos dscp-map priority` command).

The DSCP Policy Table associates an 802.1p priority with a specific ToS byte codepoint in an IPv4 packet. This enables you to set a LAN policy that operates independently of 802.1Q VLAN-tagging.

In the default state, most of the 64 codepoints do not assign an 802.1p priority, as indicated by `No-override` in the table below.

You can use the following command to list the current DSCP Policy table.

Syntax:

```
show qos dscp-map
```

Displays the DSCP Policy Table.

Table 24: *Partial display from the default DSCP Policy Table*

DSCP codepoint	DSCP value	802.1p tag	DSCP policy name
000000	0		cs0
000001	1	No-override	
000010	2	No-override	
000011	3	No-override	
000100	4	No-override	
000101	5	No-override	

Table Continued

DSCP codepoint	DSCP value	802.1p tag	DSCP policy name
000110	6	No-override	
000111	7	No-override	
001000	8	1	cs1
001001	9	No-override	
001010	10	No-override	af11
001011	11	No-override	

Default priority settings for selected codepoints

A few codepoints have default priorities assigned. Some are implied by the DSCP standards for Assured-Forwarding (af) and Expedited-Forwarding (ef). In addition, Class Selector (cs) defaults are used to automap codepoints that overlap IP-Precedence mode.

You can change the priorities for the default policies by using `qos dscp-map <codepoint> priority <0-7>` . (These policies are not in effect unless you have either applied the policies to a QoS classifier or configured QoS Type-of-Service to be in `diff-services` mode.)

Quickly listing non-default codepoint settings

The DSCP Policy Table in **Differentiated Services Codepoint (DSCP) mapping**, lists the switch's default codepoint/priority settings. If you change the priority of any codepoint setting to a non-default value and then execute `write memory`, the switch will list the non-default setting in the `show config` display. For example, in the default configuration, the following codepoint settings are true:

Codepoint	Default priority
af11	No-override
13	No-override
af13	No-override

If you change all three settings to a priority of 3, and then execute `write memory`, the switch will reflect these changes in the `show config` listing:

Figure 34: *Example of show config listing with non-default priority settings in the DSCP table*

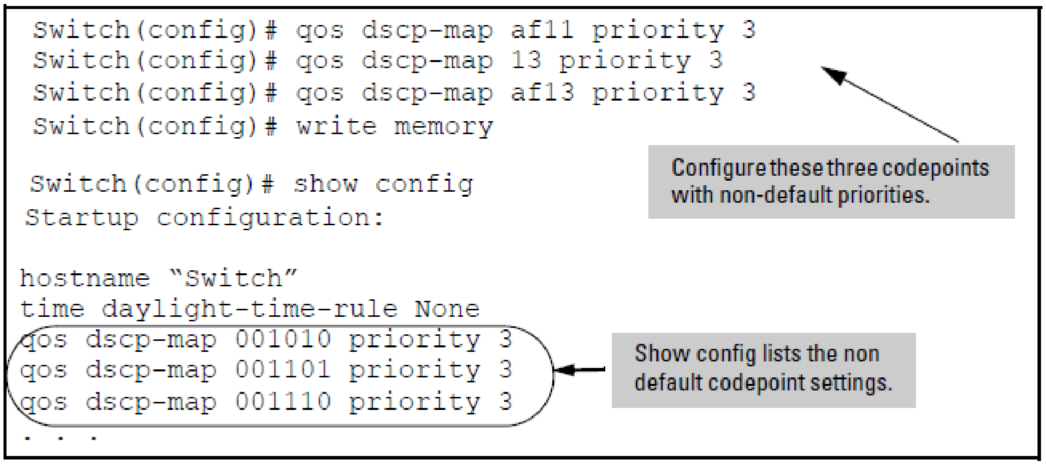

```
 Switch(config)# qos dscp-map af11 priority 3
 Switch(config)# qos dscp-map 13 priority 3
 Switch(config)# qos dscp-map af13 priority 3
 Switch(config)# write memory

 Switch(config)# show config
 Startup configuration:

hostname "Switch"
time daylight-time-rule None
qos dscp-map 001010 priority 3
qos dscp-map 001101 priority 3
qos dscp-map 001110 priority 3
. . .
```

Configure these three codepoints with non-default priorities.
Show config lists the non default codepoint settings.

Effect of `No-override`: In the QoS Type-of-Service differentiated services mode, a `No-override` assignment for the codepoint of an outbound packet means that QoS is effectively disabled for such packets. That is, QoS does not affect the packet queuing priority or VLAN tagging.

In this case, the packets are handled as follows (as long as no other QoS feature creates priority assignments for them):

802.1Q status	Outbound 802.1p priority
Received and Forwarded on a tagged port member of a VLAN.	Unchanged
Received on an Untagged port member of a VLAN; Forwarded on a tagged port member of a VLAN.	0 (zero)—”normal”
Forwarded on an Untagged port member of a VLAN.	None

Note on changing a priority setting

If a QoS classifier is using a policy (codepoint and associated priority) in the DSCP Policy table, you must delete or change this usage before you can change the priority setting on the codepoint. Otherwise the switch blocks the change and displays this message:

```
Cannot modify DSCP Policy < codepoint > - in use by other
qos rules.
```

In this case, use `show qos <classifier>` to identify the specific classifiers using the policy you want to change; that is:

```
show qos port-priority
```

```
show qos vlan-priority
```

```
show qos type-of-service
```

For example, suppose that the 000001 (dscp 1) codepoint has a priority of 6, and several classifiers use the 000001 codepoint to assign a priority to their respective types of traffic. If you wanted to change the priority of codepoint 000001, you would do the following:

1. Identify which QoS classifiers use the codepoint.
2. Change the classifier configurations by assigning them to a different DSCP policy, or to an 802.1p priority, or to `No-override`.
3. Reconfigure the desired priority for the 000001 (dscp 1) codepoint.
4. Either reassign the classifiers to the 000001 (dscp 1) codepoint policy or leave them as they were after step 2, above.

Changing the priority setting on a policy when one or more classifiers are currently using the policy (example)

Suppose that codepoint 1 is in use by one or more classifiers. If you try to change its priority, you see a result similar to the following:

Error message for changing the priority on a DSCP policy

```
switch(config)# qos dscp-map 1 priority 2
Cannot modify DSCP Policy 1 - in use by other qos rules.
```

In this case, you would use steps similar to the following to change the priority.

1. Identify which classifiers use the codepoint you want to change, using the `show qos` *`<classifier>`* command.
2. Change the classifier configurations by assigning them to a different DSCP policy, or to an 802.1p priority, or to `No-override`.
3. Reconfigure the desired priority for the 000001 (dscp 1) codepoint.

   ```
   switch(config)# qos dscp-map 1 priority 4
   ```
4. You could now re-assign the classifiers to the original policy codepoint or leave them as currently configured.

IP Multicast (IGMP) interaction with QoS

IGMP high-priority-forward causes the switch to service the subscribed IP multicast group traffic at high priority, even if QoS on the switch has relegated the traffic to a lower priority. This does not affect any QoS priority settings, so the QoS priority is honored by downstream devices. However, QoS does take precedence over IGMP normal-priority traffic.

The switch's ability to prioritize IGMP traffic for either a normal or high priority outbound queue overrides any QoS criteria, and does not affect any 802.1p priority settings the switch may assign. For a given packet, if both IGMP high priority and QoS are configured, the QoS classification occurs and the switch marks the packet for downstream devices, but the packet is serviced by the high-priority queue when leaving the switch.

IGMP high priority	QoS configuration affects packet	Switch port output queue	Outbound 802.1p setting (requires tagged VLAN)
Not Enabled	Yes	Determined by QoS	Determined by QoS
Enabled	See above paragraph.	High	As determined by QoS if QoS is active.

Outbound queue monitor

NOTE: This feature is not supported on switches J9779A, J9780A, J9782A, and J9783A.

When QoS is used to prioritize traffic, different kinds of traffic can be assigned to different egress queues. If there is a great deal of traffic, it is desirable to be able determine if some traffic to the lower priority queues was dropped. This feature allows the egress queues for one port to be monitored for dropped packets.

Syntax:

```
qos watch-queue <port> out

no qos watch-queue <port> out
```

Configures the switch to start monitoring the specified port for the dropped packets for each queue. Disabling and then re-enabling monitoring on a port clears the per-queue dropped packet counters. For example:

```
switch(config)#: qos watch-queue 5 out
```

The `no` form of the command stops the collection of dropped traffic information.

Default: disabled.

Displaying per-queue counts

The `show interface queues` command displays the number of dropped packets for each queue for the configured port. The port must have been configured with the `qos watch-queue` command. Ports that have not been configured display zero values for the queue counts.

Monitoring egress queues on a port

```
switch(config)#: show interface queues 5

Status and Counters - Queue Counters for port 5

  Name  :
  MAC Address      : 001c2e-95ab3f
  Link Status      : Up
  Port Totals (Since boot or last clear) :
   Rx Ucast Pkts    : 142,181               Tx Ucast Pkts   : 552
   Rx B/Mcast Pkts : 10,721,488             Tx B/Mcast Pkts : 11,765
   Rx Bytes         : 1,267,216,218         Tx Bytes        : 2,652,372
   Rx Drop Packets : 0                      Tx Drop Packets : 0
  Egress Queue Totals (Since boot or last clear) :
   Queue CoS  Dropped Packets
   1     1-2  123456789012345
   2     0,3  12345678
   3     4-5  1234
   4     6-7  0
```

Configuring traffic templates

In order to define mappings of 802.1p priorities to queues, traffic class groups (traffic groups) are defined as part of a traffic template. A traffic group consists of a set of 802.1p priorities that are mapped to the same egress queue. A traffic template consists of a set of traffic groups that cover all priorities 0- 7. The number of traffic groups allowed within a traffic template is equal to the number of queues configured on a platform, although a queue may have no priorities mapped to it.

For example, if you want to configure a switch with a priority to queue mapping that matches a connected networking device's queue configuration, and the neighbor device has 3 queues configured, all priorities should be mapped to the 3 highest priority queues on the switch in a manner that matches the configuration of the neighbor device.

Once defined, the traffic template is then assigned in its entirety to all ports of the switch.

Two predefined traffic templates are provided that correspond to the IEEE 802.1p traffic group specification (default-tcgt) and the 802.1q update (dot1q.tcgt). These predefined templates may be applied as is, or they may be used as the basis for defining a custom template.

Displaying traffic template information

To display a summary of all traffic templates, enter the `show qos traffic-template` command.

List of the currently defined traffic templates and their status

```
switch(config)# show qos traffic-template

System default template: default-tcgt

Template Name                        Status      Queues
-------------                        ------      ------
example                              Valid          4
dot1q-tcgt (predefined)              Valid          4
default-tcgt (predefined)            Active         4
```

To display detailed information about a single traffic template, enter the `show qos traffic-template <template-name>` command.

List of the currently defined traffic templates and their status

```
switch(config)# show qos traffic-template default-tcgt

  Template Name: default-tcgt (predefined)
  Status       : Active

Queue-no  Priorities     Name
--------  -------------  --------------------
1         1,2            background-tcg
2         0,3            best-effort-tcg
3         4,5            controlled-load-tcg
4         6,7            control-tcg
```

Note: The Name column shows the descriptive names assigned to the traffic group to help identify their purpose.

Creating a traffic template

A traffic template can be created, modified, or deleted by entering this command in the global configuration context.

Syntax:

```
qos traffic-template <template-name>

no qos traffic-template <template-name>
```

Enter the Traffic Template Configuration level or execute one command for that level. With no additional parameters specified, the `qos traffic-template` command changes the context to Traffic Template Configuration level for execution of configuration changes to the named template.

If *<template-name>* does not exist already, the template will be created as a copy of default-tcgt and the traffic template context is entered. The maximum length is 40 characters.

If *<template-name>* already exists, the traffic template context for that template is entered and you can make modifications to the template.

```
qos traffic-template <template-name> <copy-from-name>
```

When executed, a new template is created with the name *<template-name>* as a copy of the existing template named *<copy-from-name>*.

NOTE: You cannot modify any predefined traffic templates.

```
no qos traffic-template <template-name>
```

Deletes the specified traffic template. Predefined templates and the currently active template may not be deleted.

```
show qos traffic-template [<template-name>]
```

When entered without the name of a specific traffic template, displays a list of the current traffic templates and their status - active, valid, or invalid.

When executed with the name of a specific traffic template, detailed information for that traffic template is displayed.

After executing the `qos traffic-template` *<template-name>* command in the global configuration context, you are in traffic template context and can begin modifying a newly created template or an existing template.

Creating a new traffic template and entering traffic template configuration context

```
switch(config)# qos traffic-template example

switch(cfg-tcgt-example[1])# show qos traffic-template example

 Template Name: example
 Status : Valid

Queue-no  Priorities     Name
--------  -------------  --------------------
1         1,2            background-tcg
2         0,3            best-effort-tcg
3         4,5            controlled-load-tcg
4         6,7            control-tcg
```

[1] Traffic template context

You should always check the traffic template status using the `show qos traffic-template` command to make sure that the template is valid. Invalid templates cannot be applied to ports. A traffic template will be invalid if the priorities are not mapped to an egress queue. This can happen, for example, if you delete a priority from a queue and do not reassign it to another queue.

If a template has been applied, it is considered to be in use and cannot be deleted. The `show qos queue-config` command displays the currently active traffic template.

Displaying the currently active traffic template

```
switch(config)# show qos queue-config

 Egress Queue Configuration

 Number of Queues: 8
 Traffic Template: default-tcgt

        802.1p
 Queue  Priority
 -----  --------
 1      1
 2      2
 3      0
 4      3
 5      4
 6      5
 7      6
 8      7
```

Configuring traffic groups within a traffic template

When in the traffic template context, you can modify traffic groups within that template by changing which priorities are assigned to what queues, as well as assigning a name to each group.

Syntax:

```
map-traffic-group <queue-num> priority <priority> | [name <namestring>]

no map-traffic-group <queue-num> priority <priority> | [name <namestring>]
```

Allows configuration of traffic groups with a traffic template.

The `no` form of the command removes a priority from the currently mapped queue.

`<queue-num>`: Ranges from 1 to the number of active queues, which can vary from 1 to 8 queues. The number of queues is configured with the `qos queue-config` command.

`priority <0-7>`: Specifies an 802.1p priority to assign this queue. This command may be repeated to assign multiple priorities to the same queue.

`name` *`<namestring>`*: Assigns a documentary label to the traffic group. Maximum length is 40 characters.

The `no` form of the command removes the descriptive name from the group.

Moving a priority from one traffic group to another

Typically, when modifying a traffic template, priorities are reassigned to different queues. When moving a priority from one traffic group to another, you must first delete that priority from the traffic group it is currently mapped to (`no map-traffic-group <queue-num> priority <priority>` and then add it the desired traffic group.

NOTE: Before changing the traffic groups' priority mappings from the defaults, it is important to examine the current policies that act on 802.1p priorities to ensure that the policies best serve the behavior desired on the network, including DSCP, VLAN, interface, or protocol rules.

Removing priorities from a traffic group

```
switch(cfg-tcgt-example)# no map-traffic-group 2 priority 0
```

After removing the priority from the currently mapped queue, the template becomes invalid because priority 0 is no longer mapped to any queue.

Mapping the priority to a new queue results in a valid traffic template again.

Invalid traffic template because a priority is unmapped

```
switch(cfg-tcgt-example)# show qos traffic-template example

TRAFFIC-TEMPLATE: example
Status          : Invalid

Queue-no  Priorities     Name
--------  -------------  --------------------
1         1,2            background-tcg
2         3              best-effort-tcg
3         4,5            controlled-load-tcg
4         6,7            control-tcg
```

Valid traffic template with remapped queue

```
switch(cfg-tcgt-example)# map-traffic-group 1 priority 0

switch(cfg-tcgt-example)# show qos traffic-template example

Template Name: example
Status       : Valid

Queue     Priorities     Name
--------  -------------  --------------------
1         1,2            background-tcg
2         3,0            best-effort-tcg
3         4,5            controlled-load-tcg
4         6,7            control-tcg
```

After modifying a traffic template, you must apply it to the switch to activate the new mapping. See **Applying a traffic template** on page 208.

Applying a traffic template

After creating a traffic template with the desired queue assignments, you must apply it. The same traffic templates is applied to the all ports on the switch. A reboot is required for the new template to take effect.

Syntax:

```
qos queue-config traffic-template <template-name>
```

Applies the specified traffic template to all the ports and reboots the switch.

Applying a traffic template to a switch

```
switch(config)# qos queue-config traffic-template example
 This command will modify the current running configuration,
 execute 'write memory' to replace the startup configuration,
 and then reboot.

Egress queues will be configured as follows:
 4-queues
 Queue servicing: weighted-round-robin
 simple-red disabled
 Traffic template: example

Do you want to save current configuration [y/n/^C]? y
Device will be rebooted, do you want to continue [y/n]? y
```

Port QoS Trust Mode

The Port QoS Trust feature restricts which packet QoS information may be used to determine inbound queue servicing and any priority information to be permitted into the local hop.

Port QoS Trust Mode configuration allows preservation or removal of the inbound QoS priorities carried in Layer 2 (the VLAN cos or Priority CodePoint (PCP) value, known as the 802.1p priority tag) and/or in Layer 3 (the IP-ToS byte, in IP-Precedence or IP-Diffserv mode). The different modes let the customer trust all, some, or no packet priority fields.

The per-port configuration enables the customer to trust some sources or devices and not others. This feature is mutually exclusive with any active port-priority configuration.

Configuration commands

qos trust

Syntax

```
qos trust [default|dot1p|dscp|ip-prec|none|device [none|<DEVICE-TYPE>]]
```

Description

Set the QoS Trust Mode configuration for the port.

Parameters

default

Trust 802.1p priority and preserve DSCP or IP-ToS.

device *<DEVICE-TYPE>*

On approved devices, trust IP-ToS Differentiated-Services in IP packets, and use the DSCP-MAP to remark the 802.1p priority. If the DSCP codepoint does not have an associated priority, the priority will be remarked to 0. On unapproved devices, trust 802.1p priority and preserve any IP- ToS values.

dot1p

Trust 802.1p priority and preserve DSCP or IP-ToS.

dscp

Trust IP-ToS Differentiated-Services in IP packets, and use the DSCP-MAP to remark the 802.1p priority. If the DSCP codepoint does not have an associated 802.1p priority, the priority will be remarked to 0.

ip-precedence

Trust IP-ToS IP-Precedence mode in IP packets and remark the 802.1p priority.

none

Do not trust either the 802.1p priority or the IP-ToS values.

QoS trust devices

aruba-ap

Aruba Access point device.

none

Clear all trusted devices from port.

NOTE:

Both SNMP and the CLI will verify that the current QoS Port Priority and desired QoS Trust Mode configuration are not mutually exclusive (and conversely).

qos dscp-map

Syntax

```
qos dscp-map <CODEPOINT> priority <PRIORITY> [name <NAME> | default | legacy]
```

Description

Modifies DSCP mapping.

Parameters

default

Returns switch to the fully mapped factory-default configuration.

legacy

Restore the legacy default behavior (partial mapping) used in earlier code releases.

Show commands

show qos trust

Syntax

```
show qos trust [device] <PORT>
```

Description

Shows port-based QoS trust configuration

Parameters

device

Show list of trusted devices per-port.

<port>

Show trusted devices on a single port.

Usage

```
show qos trust [device | [ethernet <PORT-LIST> ]
```

show qos trust

```
switch# show qos trust

Port-based qos Trust Configuration

  Port    Trust Mode    | Device Trust State ---- --- ----
  A1      Default       |
  A2      Default       |
  A3      Device**      | Trusted
  A4      IP-Prec       |
  A5      Dot1p         |
  A5      None          |
  A5      DSCP          |
  A5      Device**      |
  A5      Dot1p         |

** For a list of trusted devices per-port, use the command show qos trust device.
To show trusted devices on a single port, use the command show qos trust device <PORT>.
```

show qos trust device

```
switch# show qos trust device

Port-Based QoS Trust Configuration

  Port     Trusted Devices
  -------- ---------------
  A1       aruba-ap
  A2       aruba-ap
  A4       aruba-ap
```

show qos trust device <PORT>

```
switch# show qos trust device <PORT>

Port A4 QoS Trust Configuration
   Current state: Trusted

   Trusted Devices: aruba-ap
```

QoS queue configuration

QoS queue configuration reduces the number of outbound queues that all switch ports use to buffer packets for 802.1p user priorities. By default the switches covered in this guide use four queues. Change the default QoS queue configuration to four-queue mode or two-queue mode to increase the available bandwidth per queue.

Use the following commands to change the number of queues per port and display the current priority queue configuration on the switch.

Syntax:

```
qos queue-config <2-queues | 4-queues | queue-servicing | traffic-template>
```

Configures the number of outbound priority queues for all ports on the switch using one of the following options: 2-queues, 4-queues.

(Default: 4-queues)

CAUTION: This command will execute a `write memory` followed by an immediate reboot, replacing the Startup configuration with the content of the current Running configuration.

The new configuration will set the new number of outbound port queues.

If you select anything but 'yes' for this operation, the operation is aborted and a message stating `Operation aborted` appears.

`queue-servicing`: Set the outbound queue servicing algorithm to use on congested ports. `strict-priority`: Gives higher priority queues unconditional precedence over lower priority queues for maximum throughput of high priority traffic. This can result in lower priority traffic being completely blocked during periods of congestion. `weighted-round-robin`: Gives all outbound queues a weighted share of the port bandwidth so that lower priority queues are not completely blocked (default).

`traffic-template`: Sets the traffic template to be used on all ports to define the mapping of 802.1p packet priority values to outbound queues.

Syntax:

```
show qos queue config
```

Displays the current qos queue configuration.

Mapping of outbound port queues

This table shows the mapping of 802.1p priorities to outbound port queues:

Table 25: *Mapping 802.1p priorities to outbound port queues*

802.1p priority	4 Queues	2 Queues
1 (lowest)	1	1
2		
0 (normal)	2	
3		
4	3	2
5		
6	4	
7 (highest)		

Configuring the number of priority queues

To change the number of outbound priority queues for all ports on the switch, use the `qos queue-config` command.

CAUTION: The `qos queue-config` command executes a `write memory` followed by an immediate reboot, replacing the Startup configuration with the contents of the current Running configuration.

Example:

To change the number of outbound priority queues for all ports on the switch from four queues to two:

Procedure

1. Configure the number of outbound priority queues by using the `qos queue-config` command.

   ```
   switch(config)# qos queue-config 2-queues
   ```

2. A caution message is displayed (see the Caution note above) concluding with the following prompt:

   ```
   Do you want to save the current configuration? [y/n]

   Device will be rebooted, do you want to continue? [y/n]
   ```

3. A second confirmation prompt appears:

   ```
   Please confirm reset. [Yes/Cancel]
   ```

4. Type **‘y’** (Yes) to initiate a write memory followed by an immediate reboot (entering **‘n’** (No) at either of the two prompts will cancel the command and maintain the current queue configuration on the switch).
5. The changes will be committed to the startup configuration and the switch will reboot automatically with the new priority queue changes in effect.

Viewing the QoS queue configuration

Syntax:

```
show qos queue-config
```

Displays the current priority queue configuration per queue.

Viewing QoS queue configuration

```
switch# show qos queue-config

 Outbound Port Queue Configuration

        802.1p
 Queue Priority
 ----- --------
   1     0-3
   2     4-7
```

QoS operating notes and restrictions

- **All switches:** For explicit QoS support of IP subnets, Hewlett Packard Enterprise recommends forcing IP subnets onto separate VLANs and then configuring VLAN-based classifiers for those VLANs.
- **For devices that do not support 802.1Q VLAN-tagged ports:** For communication between these devices and the switch, connect the device to a switch port configured as `Untagged` for the VLAN in which you want the device's traffic to move.
- **Port tagging rules:** For a port on the switch to be a member of a VLAN, the port must be configured as either `Tagged` or `Untagged` for that VLAN. A port can be an untagged member of only one VLAN of a given protocol type. Otherwise, the switch cannot determine which VLAN should receive untagged traffic.
- **Not supported:** Use of an inbound 802.1p packet priority as a classifier for remapping a packet's outbound priority to different 802.1p priority. For example, where inbound packets carry an 802.1p

priority of 1, QoS cannot be configured use this priority as a classifier for changing the outbound priority to 0.

Overview of RPVST+

NOTE: For information on configuring basic and multiple instance spanning tree, see **Multiple instance spanning tree operation**.

RPVST+ is a proprietary spanning tree implementation that extends RSTP (802.1w) to run a separate spanning tree for each VLAN on the switch, and ensures that only one active, loop-free path exists between any two nodes on a given VLAN.

Table 26: *RPVST scalability*

Platform	Maximum Allowed RPVST Enabled VLANs	Recommended Maximum Virtual Ports	Maximum Allowed Virtual Ports (x is the number of logical ports in the system)
2530	32	250	250 + x

Where x is the sum of all physical ports and logical interface, such as Trk1 if configured.

The following shows how x is calculated using the 2920 switch as an example, you can substitute the actual values for your switch as shown in the table to do a similar calculation. In a stack of 4 with 2x2920-24 and 1 trunk interface Trk1 configured, x will be 24+24+1=49. Therefore the maximum allowed vPorts is 299.

```
Switch# show spanning-tree system-limits rapid-pvst

 Spanning Tree Information

  STP Enabled              : No
  Mode                     : MSTP
  RPVST Enabled VLANs      : 1,4,20-23

  Switch MAC Address                   : 40a8f0-0df69e
  Count of RPVST Enabled GVRP VLANs    : 0
  Count of RPVST Enabled VLANs         : 6
  Maximum Allowed RPVST Enabled VLANs  : 64
  Count Of Total Virtual Ports         : 51
  Maximum Allowed Virtual Ports        : 299

                       Current         Operational     Recommended Maximum
  Ports                Virtual Ports   Virtual Ports   Virtual Ports
  -------------------- --------------- --------------- --------------------
  Member 1/1-24        31              27              250
  Member 2/1-24        31              24              250
```

General steps for configuring RPVST+

The general steps for configuring RPVST+ via the CLI are:

Procedure

1. Select RPVST+ as the active spanning tree mode by entering the following command:

   ```
   spanning-tree mode rapid-pvst
   ```

   To begin with the default RPVST+ configuration (recommended), go to step 6.

2. Configure global spanning tree parameters.
3. Configure per-VLAN parameters.
4. Configure per-port per-VLAN parameters. These commands affect RPVST+ operation on traffic associated with the specified VLANs through the specified ports.
5. Configure per-port parameters. These commands affect RPVST+ operation for all traffic through the specified ports.
6. Use one of the following commands to enable RPVST+ spanning tree operation on the switch:

   a. One or more selected VLANs: `spanning-tree vlan <vid-list>`

   b. The first 32 VLANs: `spanning-tree`

      Any VLANs in excess of the first 32 would have RPVST+ disabled. In this case, use the `no spanning-tree vlan <vid-list>` command to change the mix of RPVST+ enabled and disabled VLANs.

Additional configuration options include:

- **Configuring BPDU filtering**
- **Allowing traffic on VLAN ID (PVID) mismatched links**
- **Configuring STP loop guard**

Configuring RPVST+ at a glance

The general steps for configuring RPVST+ via the CLI are:

Procedure

1. Select RPVST+ as the active spanning tree mode by entering the following command: `spanning-tree mode rapid-pvst`

   To begin with the default RPVST+ configuration (recommended), go to step 6.

2. Configure global spanning tree parameters.
3. Configure per-VLAN parameters.
4. Configure per-port per-VLAN parameters. These commands affect RPVST+ operation on traffic associated with the specified VLAN(s) through the specified port(s).
5. Configure per-port parameters. These commands affect RPVST+ operation for all traffic through the specified port(s).
6. Use one of the following commands to enable RPVST+ spanning tree operation on the switch:

   a. One or more selected VLANs:

      ```
      spanning-tree vlan vid-list
      ```

   b. One or more selected VLANs:

```
spanning-tree vlan vid-list
```

c. The first 400 VLANs:

```
spanning-tree
```

Any VLANs in excess of the first 400 would have RPVST+ disabled. In this case, use the `no` spanning-tree vlan *vid-list* command to change the mix of RPVST+ enabled and disabled VLANs.

Additional configuration options include:

- **Allowing traffic on VLAN ID (PVID) mismatched links** on page 222
- **Configuring STP loop guard** on page 223

Selecting RPVST+ as the spanning tree mode

Syntax:

```
spanning-tree mode [mstp | rapid-pvst]
```

```
no spanning-tree mode [mstp | rapid-pvst]
```

Specifies that spanning tree will run in MSTP (default) or RPVST+ mode.

To view Mode, use the `show run` command. This will eliminate confusion if there is an RPVST configuration but MSTP is running. This will lead to a change in the existing factory default setting.

RPVST+ parameters can be configured even if the mode is MSTP and vice versa. This command does not enable/disable spanning tree. It sets the mode which is operational once spanning tree is enabled using `spanning-tree enable`.

The `no` form of the command changes the spanning tree mode to the default mode (MSTP) .

Configuring global spanning tree

Syntax:

```
spanning-tree extend system-id
```

Creates a unique bridge identifier for each VLAN by adding the VLAN ID (vid) value to the priority field of the bridge identifier in every RPVST+ BPDU.

Syntax

```
no spanning-tree log state-transitions [instance <instance-id> cst]
```

- The command enables or disables event logging for port-block events.
- List of VLAN identifiers
- Range: <instance-id> 1–16

  ```
  [vlan <vid-list>]
  ```

Syntax:

```
no spanning-tree ignore-pvid-inconsistency
```

Causes the switch to ignore per-VLAN ID inconsistencies when the ports on both ends of a point-to-point link are untagged members of different VLANs, thus allowing RPVST+ to run on the mismatched links. On a given switch, affects all ports belonging to VLANs on which RPVST+ is enabled.

Default: Disabled

NOTE: The no form of this command is ineffective when there is a PVID inconsistency between a VLAN1 and any non-VLAN1 member because VLAN1 uses IEEE BPDUs to form a spanning tree topology.

Syntax:

```
no spanning-tree bpdu-protection-timeout timeout
```

- Configures the duration of time when protected ports receiving unauthorized BPDUs will remain disabled. The default value of 0 (zero) sets an infinite timeout (that is, ports that are disabled by bpdu–protection are not, by default, re-enabled automatically).
- Default: 0
- Range: 0 - 65535 seconds

Configuring per-VLAN spanning tree

Syntax:

```
spanning-tree vlan <vid-list> hello-time 1...10
```

Specifies the time in seconds between transmissions of BPDUs on the specified VLAN(s) when the switch is root for those VLAN(s).

Default: 2

Range: 1 - 10

Syntax:

```
spanning-tree vlan <vid-list>forward-delay 4...30
```

Sets the time in seconds the switch waits before transitioning from listening to learning and from learning to forwarding states.

Default: 15

Range: 4 - 30

Syntax:

```
spanning-tree vlan vid-listmaximum age 6...40
```

Sets the maximum age in seconds of received STP information before it is discarded for specified VLAN(s).

Default: 20

Range: 6 - 40

NOTE: Maximum age must be within the following bounds:

- greater than or equal to 2x (hello-time +1)
- less than or equal to 2x (forward-delay - 1)

Syntax:

```
spanning-tree vlan vid-list priority 0...15
```

Sets the switch (bridge) priority for the designated VLAN. The switch compares this priority with the priorities of other switches on the same VLAN to determine the RPVST+ root switch for the VLAN. The lower the priority value, the higher the priority. The switch with the lowest Bridge Identifier on the VLAN is elected as the RPVST+ root switch for that VLAN.

The Bridge Identifier is composed of a configurable Priority (2 bytes) and the switch's MAC address (6 bytes). You can change the Priority provides flexibility for determining which switch on the VLAN will be the root for RPVST+, regardless of its MAC address.

The priority range for an RPVST+ switch is 0-61440. However, this command specifies the priority as a multiplier (0 - 15) of 4096. That is, when you specify a priority multiplier value of 0 - 15, the actual priority assigned to the switch is: (priority-multiplier) x 4096.

For example, if you configure "2" as the priority-multiplier on a given RPVST+ switch, then the Switch Priority setting for the specified VLAN is 8,192.

NOTE: If multiple switches on the same VLAN have the same priority setting, then the switch with the lowest MAC address becomes the root switch for that VLAN.

Syntax:

```
spanning-tree vlan vid-list root {primary | secondary}
```

```
no spanning-tree vlan vid-list root {primary | secondary}
```

Specifies the switch as the primary or secondary root bridge for the specified VLAN(s). Otherwise, by default, the root bridge for each VLAN will be determined by the lowest MAC address in that topology.

The `no` form of the command returns the determination of root to the lowest MAC address criterion.

Configuring per-port per-VLAN spanning tree

Syntax

```
spanning-tree pathcostrapid-pvst | mstp[8021d | 8021t | proprietary]
```

```
no spanning-tree pathcostrapid-pvst | mstp[8021d | 8021t | proprietary]
```

Specify a standard to use when calculating the default pathcost.

Default: 8021t

NOTE: All devices in the network should be configure to use same pathcost mode for proper functioning.

Syntax:

```
spanning-tree port port-#: vlan vid-list path-cost {auto | [1...200000000]}
```

```
no spanning-tree port port-#: vlan vid-list path-cost {auto | [1...200000000]}
```

Sets the path cost for a single port on the specified VLAN(s). If the port is a member of more than one VLAN, the `path-cost` applies only where the port has traffic for the VLAN(s) specified.

Default: auto

Range: 1 - 200000000

The `no` form of the command returns `path-cost` to its default setting.

Syntax:

```
spanning-tree port port-number vlan vid-list priority <0-15> path-costauto | <Path-Cost>
```

```
no spanning-tree port port-number vlan vid-list priority <0-15> path-costauto | <Path-Cost>
```

Sets the port priority for the specified VLANs. The value is in the range of 0-240 divided into steps of 16 that are numbered 0 to 15. The default is step 16.

The per-port per-VLAN priority is used to help choose the root port for a switch on the specified VLAN if there are multiple links to the root switch.

Default: 8

Range 0 - 15

The `no` form of the command sets the priority to its default value.

Configuring per-port spanning tree

Syntax:

```
spanning-tree port-list admin-edge-port
```

```
no spanning-tree port-list admin-edge-port
```

Enables `admin-edge-port` on ports connected to end nodes. During spanning tree establishment, ports with `admin-edge-port` enabled transition immediately to the forwarding state. If a bridge or switch is detected on the segment, the port automatically operates as non-edge, not enabled.

If `admin-edge-port` is disabled on a port and `auto-edge-port` has not been disabled, the `auto-edge-port` setting controls the behavior of the port.

Default: No - disabled

The `no` form of the command disables edge-port operation on the specified ports.

Syntax:

```
spanning tree port-list auto-edge-port
```

```
no spanning tree port-list auto-edge-port
```

Enables or disables the automatic identification of edge ports. The port will look for BPDUs for 3 seconds. If there are none, it begins forwarding packets. If `admin-edge-port` is enabled for a port, the setting for auto-edge-port is ignored whether set to yes or no. If `admin-edge-port` is set to No, and `auto-edge- port` has not been disabled (set to No), then the auto-edge-port setting controls the behavior of the port.

Default: Yes - enabled

The `no` form of the command disables `auto-edge-port operation` on the specified ports.

Syntax:

```
no spanning tree port-list bpdu-filter
```

Enables or disables BPDU filtering on the specified port(s). The `bpdu-filter` option forces a port to always stay in the forwarding state and be excluded from standard STP operation.

Default: Disabled

Syntax:

Enables or disables BPDU protection on the specified port(s).

Syntax:

```
spanning tree port-list point-to-point-mac [true | false | auto]
```

Informs the switch of the type of device to which a specific port connects.

`true` (default)

Indicates a point-to-point link to a device such as a switch, bridge, or end-node.

false

Indicates a connection to a hub (which is a shared LAN segment).

auto

Causes the switch to set Force-False on the port if it is not running at full duplex. (Connections to hubs are half-duplex.)

Syntax:

```
spanning tree port-list root-guard
```

This feature is available in RPVST+ only. When a port is enabled as `root-guard`, it cannot be selected as the root port even if it receives superior STP BPDUs. The port is assigned an “alternate” port role and enters a blocking state if it receives superior STP BPDUs. (A superior BPDU contains “better” information on the root bridge and path cost to the root bridge, which would normally replace the current root bridge selection.)

The superior BPDUs received on a port enabled as `root-guard` are ignored. All other BPDUs are accepted and the external devices may belong to the spanning tree as long as they do not claim to be the Root device. Use the following command on RPVST+ switch ports that are connected to devices located in other administrative network domains to ensure the stability of the core RPVST+ network topology so that undesired or damaging influences external to the network do not enter.

Default: Disabled.

Syntax:

```
spanning-tree port-list tcn-guard
```

When `tcn-guard` is enabled for a port, it causes the port to stop processing or propagating received topology change notifications and topology changes to other ports.

Default: Disabled.

Enabling or disabling RPVST+ spanning tree

With the spanning tree mode set to RPVST+, you can do either of the following:

- Enable or disable RPVST+ on all VLANs on the switch.
- Enable or disable RPVST+ on specified VLANs that are RPVST+-enabled on the switch.

Syntax:

```
no spanning-tree [enable | disable]
```

To globally enable RPVST+ on all VLANs on the switch, use either of the following:

```
spanning-tree [enable]
```

```
no spanning-tree disable
```

To globally disable RPVST+ on all VLANs on the switch, use any of the following:

```
no spanning-tree
```

```
spanning-tree disable
```

```
no spanning-tree enable
```

NOTE: This status will always be shown in `show run` to let you know whether the spanning-tree is enabled. Having spanning tree present but not enabled will lead to a change in the existing factory default settings.

NOTE: This command overrides the per-VLAN enable/disable command (below).

Syntax:

```
spanning-tree vlan vid list [enable | disable]
```

To enable RPVST+ on one or more VLANs on the switch, use either of the following:

```
spanning-tree vlan vid list enable
```

```
no spanning-tree vlan vid list disable
```

To disable RPVST+ on one or more VLANs on the switch, use any of the following:

```
no spanning-tree vlan vid list
```

```
spanning-tree vlan vid listdisable
```

```
no spanning-tree vlan vid list enable
```

Allowing traffic on VLAN ID (PVID) mismatched links

When RPVST+ is running in the default configuration on a link where there is a VLAN ID mismatch, PVST blocks the link, resulting in traffic on the mismatched VLANs being dropped. However, there can be instances where traffic passing between mismatched VLANs on a link is desirable. When enabled on the switch, the `ignore-pvid-inconsistency` command allows this behavior. That is, where the ports on both ends of a point-to-point link are untagged members of different VLANs, enabling `ignore-pvid-inconsistency` enables RPVST+ to process untagged RPVST+ BPDUs belonging to the peer's untagged VLAN as if it was received on the current device's untagged VLAN.

Syntax:

```
no spanning-tree ignore-pvid-inconsistency
```

Causes the switch to ignore per-VLAN ID inconsistencies when the ports on both ends of a point-to-point link are untagged members of different VLANs, thus allowing RPVST+ to run on the mismatched links. On a given switch, this affects all ports belonging to VLANs on which RPVST+ is enabled.

Default: Disabled

RPVST+ behavior

Table 27: *RPVST+ behavior with ignore-pvid-inconsistency enabled*

Switch “A” Port on VLAN X	Switch “B” Peer port on VLAN Y	RPVST+ behavior with ignore-pvid-inconsistency enabled
Untagged on VLAN 10	Untagged on VLAN 10	Forward
Untagged on VLAN 10	Untagged on VLAN 20	Forward[11,]
Untagged on VLAN X	Tagged on VLAN X	Drop
Untagged on VLAN X	Tagged on VLAN Y	Drop (traffic from both VLANs)
Tagged on VLAN X	Tagged on VLAN X	Forward[1]
Tagged on VLAN X	Tagged on VLAN Y	Drop (traffic from both VLANs)

[1] Forwarding state applies if the link has not been blocked by RPVST+ as a redundant link.

NOTE: The `no spanning-tree ignore-pvid-inconsistency` command is ineffective when there is a PVID inconsistency between a VLAN1 and any non-VLAN1 member because VLAN1 uses IEEE BPDUs to form a spanning tree topology.

Configuring STP loop guard

Spanning tree is used to ensure a loop-free topology over the LAN. Occasionally a hardware or software failure can cause STP to fail, creating STP/ forwarding loops that can cause network failures where unidirectional links are used. The non-designated port transitions in a faulty manner because the port is no longer receiving STP BPDUs.

STP Loop Guard causes the non-designated port to go into the STP loop inconsistent state instead of the forwarding state. In the loop-inconsistent state, the port prevents data traffic through the link, therefore avoiding the loop creation. When BPDUs again are received on the inconsistent port, it resumes normal STP operation automatically.

Syntax:

```
spanning-tree port-list loop-guard
```

```
no spanning-tree port-list loop-guard
```

Enables STP Loop Guard on a particular port or ports. STP Loop Guard is best applied on blocking or forwarding ports.

The `no` form of the command disables STP Loop Guard.

Default: Disabled

Figure 35: *Loop creation with transmission failure*

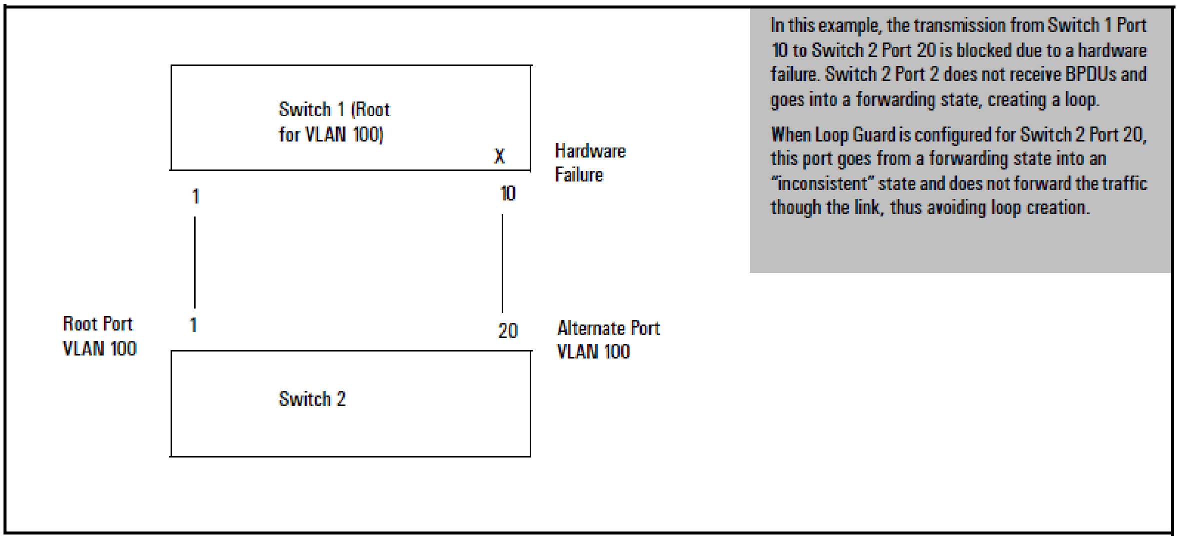

In this example, the transmission from Switch 1 Port 10 to Switch 2 Port 20 is blocked due to a hardware failure. Switch 2 Port 2 does not receive BPDUs and goes into a forwarding state, creating a loop.

When Loop Guard is configured for Switch 2 Port 20, this port goes from a forwarding state into an "inconsistent" state and does not forward the traffic though the link, thus avoiding loop creation.

Before configuring loop guard

Before configuring Loop Guard on port 20, the status of VLAN 20 appears as follows:

```
switch(config)# show spanning-tree vlan 20

 Spanning Tree Information

  STP Enabled              [No] : Yes
  Mode                          : RPVST
  Extended System ID            : Enabled
  Ignore PVID Inconsistency     : Disabled
  Switch MAC Address            : 002347-c651c0


  VLAN ID                       : 20
  RPVST Enabled                 : Enabled


  Root MAC Address              : 0024a8-d13a40
  Root Priority                 : 32,768
  Root Path Cost                : 20,000
  Root Port                     : 1
  Operational Hello Time (secs) : 2
  Topology Change Count         : 2
  Time Since Last Change        : 9 secs

                                                           Designated
  Port   Type       Cost   Priority  Role       State      Bridge
  -----  ---------- ------ --------- ---------- ---------- ---------------
   1     100/1000T  20000  128       Root       Forwarding 0024a8-d13a40
  20     10/100TX   200000 128       Alternate  Blocking   002347-587b80
```

After configuring loop guard

By executing `spanning-tree 20 loop-guard`, loop guard has been configured on port 20 of Switch 2:

```
switch(config)# show spanning-tree

 Spanning Tree Information

  STP Enabled           [No] : Yes
  Mode                       : RPVST
  Extended System ID         : Enabled
  Ignore PVID Inconsistency  : Disabled
  RPVST Enabled VLANs        : 20


  Switch MAC Address         : 002347-c651c0
  Root Guard Ports           :
  Loop Guard Ports           : 20
  TCN Guard Ports            :
  BPDU Protected Ports       :
  BPDU Filtered Ports        :
  Auto Edge Ports            : 1-24
  Admin Edge Ports           :

  VLAN  Root Mac        Root       Root       Root                  Hello
  ID    Address         Priority   Path-Cost  Port                  Time(sec)
  ----- --------------- ---------- ---------- --------------------- ---------
  100   0024a8-d13a40   32,768     20,000     1                     2
```

Switch ceasing to send BPDUs

With switch 1 ceasing to send BPDUs through port 20 to switch 2, port 20 goes into the “inconsistent” state and ceases to forward traffic, as displayed in the following `show spanning-tree` output for VLAN 20.

```
switch(config)# show spanning-tree vlan 20

 Spanning Tree Information

  STP Enabled               [No] : Yes
  Mode                           : RPVST
  Extended System ID             : Enabled
  Ignore PVID Inconsistency      : Disabled
  Switch MAC Address             : 002347-c651c0


  VLAN ID                        : 20
  RPVST Enabled                  : Enabled


  Root MAC Address               : 0024a8-d13a40
  Root Priority                  : 32,768
  Root Path Cost                 : 20,000
  Root Port                      : 1
  Operational Hello Time (secs)  : 2
  Topology Change Count          : 3
  Time Since Last Change         : 42 hours

                                                          Designated
  Port  Type       Cost   Priority  Role       State      Bridge
  ----- ---------- ------ --------- ---------- ---------- ---------------
  1     100/1000T  20000  128       Root       Forwarding 0024a8-d13a40
  20    10/100TX   200000 128       Alternate  Inconsi... 002347-587b80
```

Viewing configuration file with loop guard enabled

The following example displays `show spanning-tree config` output with loop guard enabled on Port 20:

```
switch(config)# show spanning-tree config

 Spanning Tree Information

  STP Enabled           [No] : Yes
  Mode                       : RPVST
  Extended System ID         : Enabled
  Ignore PVID Inconsistency : Disabled
  RPVST Enabled VLANs        : 100


  Switch MAC Address         : 002347-c651c0

  Root Guard Ports           :
  Loop Guard Ports           : 20
  TCN Guard Ports            :
  BPDU Protected Ports       :
  BPDU Filtered Ports        :
  Auto Edge Ports            : 1-24
  Admin Edge Ports           :

                Max Age Forward    Hello     Admin Root
  VLAN Priority (sec)   Delay(sec) Time(sec) Bridge
  ---- -------- ------- ---------- --------- ----------------
  100  32768    20      15         2         Not Configured
```

About RPVST+

RPVST+ is a proprietary spanning tree implementation that extends RSTP (802.1w) to run a separate spanning tree for each VLAN on the switch, and ensures that only one active, loop-free path exists between any two nodes on a given VLAN.

Comparing spanning tree options

Without spanning tree, having more than one active path between a pair of nodes causes loops in the network, which can result in duplication of messages, leading to a “broadcast storm” that can bring down the network.

The 802.1D spanning tree protocol operates without regard to a network's VLAN configuration, and maintains one common spanning tree throughout a bridged network. This protocol maps one loop-free, logical topology on a given physical topology. This results in the least optimal link utilization and longest convergence times.

The 802.1s multiple spanning tree protocol (MSTP) uses multiple spanning tree instances with separate forwarding topologies. Each instance is composed of one or more VLANs, which significantly improves network link utilization and the speed of reconvergence after a failure in the network’s physical topology. However, MSTP requires more configuration overhead and is more susceptible to dropped traffic due to misconfiguration.

Rapid spanning tree protocol (RSTP) requires less configuration overhead, provides faster convergence on point-to-point links, and speedier failure recovery with predetermined, alternate paths. The switches covered by this guide, use the IEEE Rapid Per-VLAN spanning tree Protocol (RPVST) standard. RPVST was introduced as an enhancement to Rapid spanning tree Protocol (RSTP) to improve the link utilization issue and require less configuration overhead. Basically, RPVST+ is RSTP operating per-VLAN in a single layer 2 domain. VLAN tagging is applied to the ports in a multi-VLAN network to enable blocking of redundant links in one VLAN while allowing forwarding over the same links for non-redundant use by another VLAN. Each

RPVST+ tree can have a different root switch and therefore can span through different links. Since different VLAN traffic can take different active paths from multiple possible topologies, overall network utilization increases.

Another major advantage of RPVST+ is that it localizes topology change propagation to individual VLANs. Since there is a separate spanning tree for each VLAN, topology changes affecting a particular VLAN are propagated only inside that VLAN. The switch flushes the MAC addresses learned only on the affected VLAN and other VLAN traffic is not disturbed. This minimizes the network flooding caused by the spanning tree topology changes. This is a significant improvement in the case of a large, flat, layer 2 network. In a network having a large number of per-VLAN spanning tree instances, RPVST+ can result in an increased load on the switch's CPU.

Understanding how RPVST+ operates

RPVST+ applies one RSTP tree per-VLAN. Each of these RSTP trees can have a different root switch and span the network through shared or different links. As shown in the following diagram, since the active paths for traffic on different VLANs can use the same for different links, multiple topologies are possible, and overall network utilization increases.

Figure 36: *RSTP forming a single spanning tree across all VLANs*

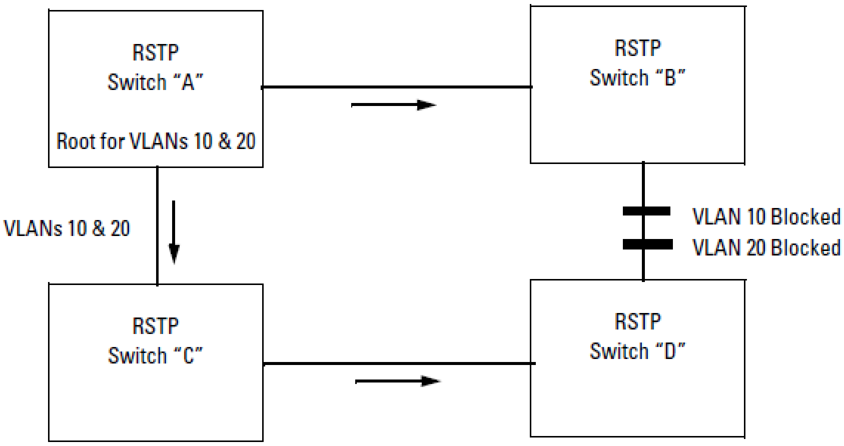

The topology has four switches running RSTP. Switch "A" is the root switch. To prevent a loop, RSTP blocks the link between switch "B" and switch "D". There are two VLANs in this network (VLAN 10 and VLAN 20). Since RSTP does not have VLAN intelligence, it forces all VLANs in a layer 2 domain to follow the same spanning tree. There will not be any traffic through the link between switch "B" and switch "D" and hence the link bandwidth gets wasted. On the other hand, RPVST+ runs different spanning trees for different VLANs. Consider the following diagrams.

Figure 37: *RPVST+ creating a spanning tree for VLAN 10*

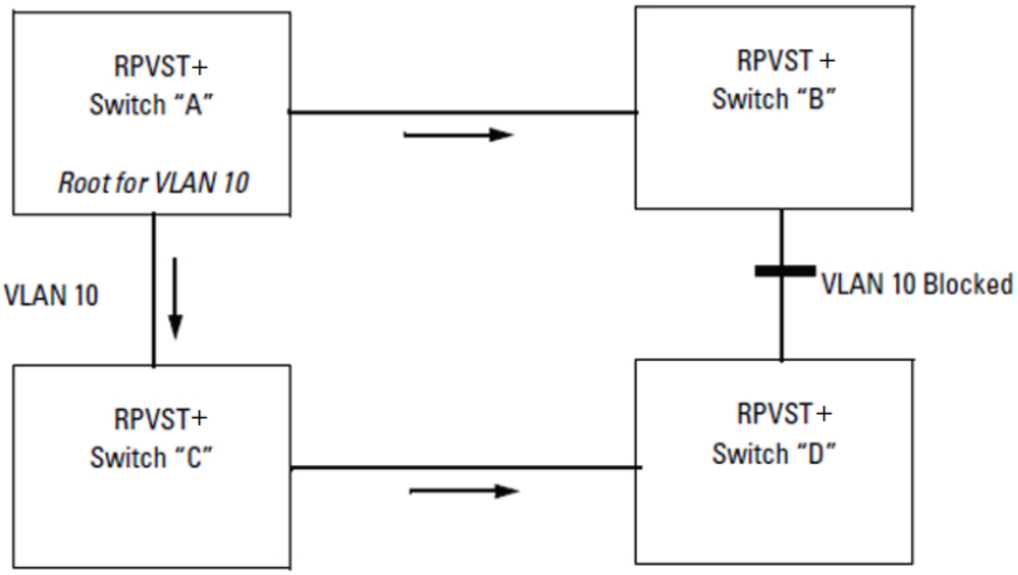

Figure 38: *RPVST+ creating a spanning tree for VLAN 20*

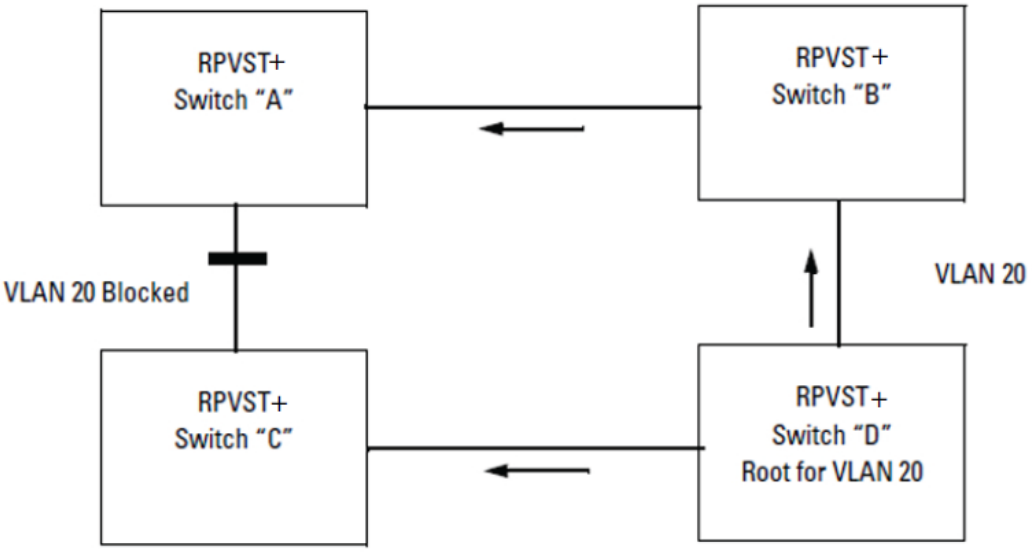

The two topologies above are the same as the first topology, but now the switches run RPVST+ and can span different trees for different VLANs. Switch "A" is the root switch for the VLAN 10 spanning tree and switch "D" is the root switch for the VLAN 20 spanning tree. The link between switch "B" and switch "D" is only blocked for VLAN 10 traffic but VLAN 20 traffic goes through that link. Similarly the link between switch "A" and switch "C" is blocked only for VLAN 20 traffic but VLAN 10 traffic goes through that link. Here, traffic passes through all the available links, and network availability and bandwidth utilization increase.

Another major advantage of RPVST+ is that it localizes topology change propagation. Since there is a separate spanning tree for each VLAN, topology changes affecting a particular VLAN are propagated only inside that VLAN. The switch flushes the MAC addresses learned only on the affected VLAN, the traffic on other VLANs is not disturbed. This minimizes the network flooding due to spanning tree topology changes. This is a significant improvement in the case of a large, flat, layer 2 network.

The following figure shows a further example of shared links and redundant path-blocking in a network running RPVST+.

Figure 39: *Sample RPVST+ network*

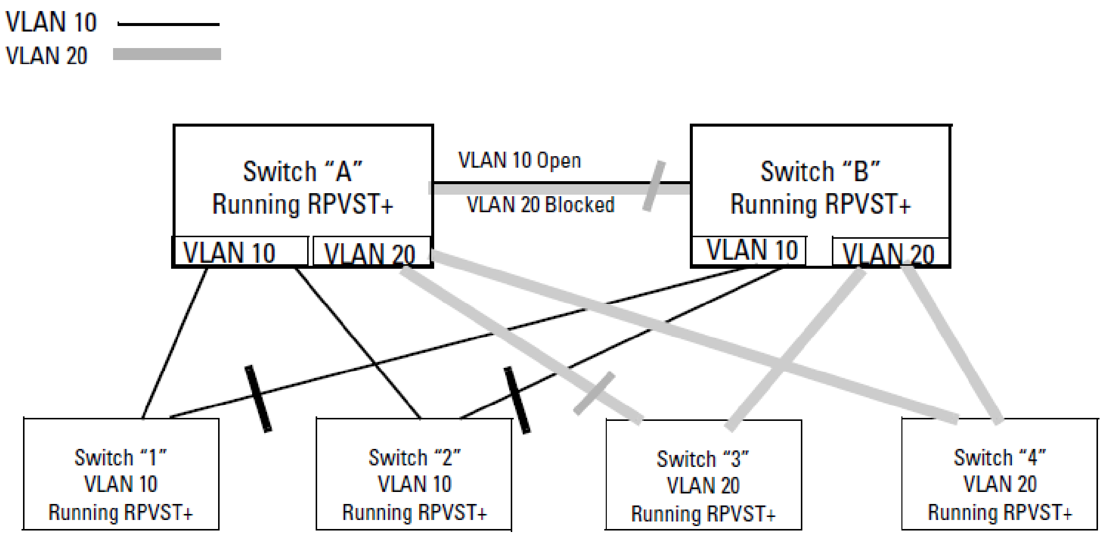

Working with the default RPVST+ configuration

In the factory default configuration, spanning tree operation is disabled. Configuring the spanning tree mode as RPVST+ on a switch and then enabling spanning tree automatically creates a spanning tree instance for each VLAN on the switch. Configuration with default settings is automatic, and in many cases does not require any adjustments. This includes operation with spanning tree regions in your network running STP, MSTP, or RSTP. Also, the switch retains its currently configured spanning tree parameter settings when spanning tree is disabled. Thus, if you disable, then later re-enable spanning tree, the parameter settings will be the same as before spanning tree was disabled.

CAUTION:

The switch automatically senses port identity and type, and automatically defines spanning tree parameters for each type, and parameters that apply across the switch. Although these parameters can be adjusted, HPE strongly recommends leaving these settings in their default configurations unless the proposed changes have been supplied by an experienced network administrator who has a strong understanding of RPVST+ operation.

RPVST+ operating notes

Recommended application

RPVST+ is ideal in networks having less than 100 VLANs. In networks having 100 or more VLANs, MSTP is the recommended spanning tree choice due to the increased load on the switch CPU.

VLAN membership

A port will be part of a given VLAN spanning tree only if the port is a member of that VLAN.

RPVST+ interoperates with RSTP and MSTP on VLAN 1

Because a switch running RPVST+ transmits IEEE spanning tree BPDUs, it can interoperate with IEEE RSTP and MSTP spanning tree regions, and opens or blocks links from these regions as needed to maintain a loop-free topology with one physical path between regions.

NOTE: RPVST+ interoperates with RSTP and MSTP only on VLAN 1.

Single spanning tree applications

One spanning tree variant can be run on the switch at any given time. On a switch running RPVST+, MSTP cannot be enabled. However, any MSTP-specific configuration settings in the startup configuration file will be maintained.

Exclusions

The following features cannot run concurrently with RPVST+:

- Features that dynamically assign ports to VLANs:
  - GVRP
  - RADIUS-based VLAN assignments (802.1X, WebAuth, MKAC auth
  - Auth-VID/UnAuth-VID configuration on interfaces
  - MAC-Based VLANs
  - LLDP Radio Port VLAN
- Switch Meshing
- QinQ
- Protocol VLANs
- Distributed Trunking
- Filter Multicast in rapid-PVST mode (The multicast MAC address value cannot be set to the PVST MAC address 01:00:0c:cc:cc:cd.)

GVRP

Spanning tree mode cannot be set to RPVST+ when GVRP is enabled, and GVRP cannot be enabled when RPVST+ is enabled.

RPVST+ operating limits

Virtual ports (vPorts) on a switch are determined by the number of physical ports on the switch, plus other factors. Exceeding the recommended number of vPorts can result in dropped BPDUs.

Allowing traffic on per-VLAN ID (PVID) mismatched links

The switch generates an Event Log message for a VID mismatch on an active RPVST+ VLAN only if `ignore-pvid-inconsistency` is disabled (the default).

If `ignore-pvid-inconsistency` is enabled on multiple switches connected by hubs, there could be more than two VLANs involved in PVID mismatches that will be ignored by RPVST+.

If there is an actual misconfiguration of port VLAN memberships in a network, then enabling ignore-pvid-inconsistency prevents RPVST+ from detecting the problem. This could result in packet duplication in the network because RPVST+ would not converge correctly.

Viewing RPVST+ statistics and configuration

NOTE:

RPVST+ is a superset of the STP/802.1D and RSTP/802.1w protocols, and uses the RPVST+ MIB (hpicfRpvst).

Viewing global and VLAN spanning tree status

Syntax:

```
show spanning-tree
```

Displays the switch's global and VLAN spanning tree status.

Viewing the switch's global and VLAN spanning tree status

```
switch# show spanning-tree

 Spanning Tree Information

  STP Enabled          [No] : Yes
  Mode                      : RPVST
  Extended System ID        : Disabled
  Ignore PVID Inconsistency : Disabled
  RPVST Enabled VLANs       : 10,20


  Switch MAC Address        : 0024a8-d13a40
  Root Guard Ports          :
  Loop Guard Ports          :
  TCN Guard Ports           :
  BPDU Protected Ports      : 23-24
  BPDU Filtered Ports       : 23-24
  Auto Edge Ports           : 1-24,A1-A4
  Admin Edge Ports          :

  VLAN  Root Mac        Root       Root       Root                 Hello
  ID    Address         Priority   Path-Cost  Port                 Time(sec)
  ----- --------------- ---------- ---------- -------------------- ---------
  10    0024a8-d13a40   32,768     0          This switch is root  2
  20    0024a8-d13a40   32,768     0          This switch is root  2
```

Viewing status for a specific VLAN

Syntax:

```
show spanning-tree vlan vlan-id
```

Displays detailed spanning tree information for the VLAN and the ports belonging to the specified VLAN.

Viewing status for a specific VLAN

```
switch(config)# show spanning-tree vlan 20

 Spanning Tree Information

  STP Enabled              [No] : Yes
  Mode                           : RPVST
  Extended System ID             : Disabled
  Ignore PVID Inconsistency      : Disabled
  Switch MAC Address             : 0024a8-d13a40


  VLAN ID                        : 20
  RPVST Enabled                  : Enabled


  Root MAC Address               : 0024a8-d13a40
  Root Priority                  : 32,768
  Root Path Cost                 : 0
  Root Port                      : This switch is root
  Operational Hello Time (secs)  : 2
  Topology Change Count          : 38
  Time Since Last Change         : 23 hours

                                                        Designated
  Port  Type       Cost   Priority  Role       State      Bridge
  ----- ---------- ------ --------- ---------- ---------- ---------------
  9     100/1000T  20000  128       Designated Forwarding 0024a8-d13a40
  21    100/1000T  20000  128       Designated Forwarding 0024a8-d13a40
  22    100/1000T  20000  128       Designated Forwarding 0024a8-d13a40
  23    100/1000T  200000 128       Designated Forwarding 0024a8-d13a40
  24    100/1000T  0      128                  Disabled
```

Viewing status for a specific port list

Syntax:

`show spanning-tree` *`port-list`*

Displays the spanning tree status for the designated port(s). You can list data for a series of ports and port trunks by specifying the first and last port or trunk of any consecutive series of ports and trunks. For example, to display data for port 20-24 and trk1, you would use this command: `show spanning-tree 20-42,trk1`

Viewing status for a specific port list

```
switch# show spanning-tree 22

 Spanning Tree Information

  STP Enabled     [No] : Yes
  Mode                 : RPVST
  RPVST Enabled VLANs : 10,20


  Switch MAC Address   : 0024a8-d13a40

  Port                    : 22
  Status                  : Up           Port Type            : 100/1000T
  BPDU Protection         : No           BPDU Filtering       : No
  Root Guard              : No           TCN Guard            : No
  Loop Guard              : No           Admin Edge Port      : No
  Admin PointToPoint MAC  : Yes

  VLAN   Port       Port      Port       Designated    Hello Oper  Oper
  ID     Path-Cost  Priority  State      Bridge        Time  Edge  PtP
  ------ ---------- --------- ---------- ------------- ----- ----- ------
  20     20000      128       Forwarding 0024a8-d13a40 2     No    Yes
  25     200000     128       Forwarding 002347-587b80 2     Yes   Yes
```

Viewing status per-port per-VLAN

Syntax:

`show spanning-tree` *`port-list`* `vlan` *`vlan-id`*

Displays detailed information for ports in the `port-list` in the given VLAN. This command further filters the output for `show spanning-tree` *`port-list`*.

Viewing status per-port per-VLAN

```
switch# show spanning-tree 22 vlan 20

 Spanning Tree Information

  STP Enabled    [No] : Yes
  Mode                : RPVST
  RPVST Enabled VLANs : 10,20


  Switch MAC Address  : 0024a8-d13a40

  Port                   : 22
  Status                 : Up           Port Type           : 100/1000T
  BPDU Protection        : No           BPDU Filtering      : No
  Root Guard             : No           TCN Guard           : No
  Loop Guard             : No           Admin Edge Port     : No
  Admin PointToPoint MAC : Yes

  VLAN   Port       Port      Port       Designated    Hello Oper  Oper
  ID     Path-Cost  Priority  State      Bridge        Time  Edge  PtP
  ------ ---------- --------- ---------- ------------- ----- ----- ------
  20     20000      128       Forwarding 0024a8-d13a40 2     No    Yes
```

Viewing the global RPVST+ configuration

Syntax:

`show spanning-tree config`

Displays the switch’s basic and per-VLAN spanning tree configuration.

The upper part of the output shows the switch’s global spanning tree configuration. The port listing shows the spanning tree port parameter settings for the spanning tree region operation (configured by the `spanning-tree` *`port-list`* command).

Viewing the global RPVST+ configuration

```
switch# show spanning-tree config

 Spanning Tree Information

  STP Enabled          [No] : Yes
  Mode                      : RPVST
  Extended System ID        : Enabled
  Ignore PVID Inconsistency : Disabled
  RPVST Enabled VLANs        : 10,20


  Switch MAC Address         : 002347-587b80

  Root Guard Ports           :
  Loop Guard Ports           :
  TCN Guard Ports            :
  BPDU Protected Ports       :
  BPDU Filtered Ports        :
  Auto Edge Ports            : 1-24
  Admin Edge Ports           :

                Max Age Forward     Hello     Admin Root
  VLAN Priority (sec)   Delay(sec) Time(sec) Bridge
  ---- -------- ------- ---------- --------- ----------------
  1    32768    20      15         2         Not Configured
  10   32768    20      15         2         Not Configured
  20   32768    20      15         2         Not Configured
```

Viewing the global RPVST+ configuration per port

Syntax:

```
show spanning-tree [ethernet]port-list config
```

Lists the spanning tree port parameter settings (global and per VLAN) for only the specified port(s) and trunk(s). You can list data for a series of ports and port trunks by specifying the first and last port or trunk of any consecutive series of ports and trunks. For example, to display data for ports 9, 11, 12, 21 and trk1, use this command: `show spanning-tree 9,11,12,21,trk1 config`

Viewing the global RPVST+ configuration per port

```
switch# show spanning-tree 9,11,12,21,22 2 trk1 config

 Spanning Tree Information

  STP Enabled     [No] : Yes
  Mode                 : RPVST
  Switch MAC Address   : 002347-587b80
  RPVST Enabled VLANs  : 10,20


        Admin Auto Admin Root Loop TCN BPDU BPDU
  Port  Edge  Edge PtP   Grd  Grd  Grd Flt  Guard
  ----- ----- ---- ----- ---- ---- --- ---- -----
  9     No    Yes  True  No   No   No  No   No
  11    No    Yes  True  No   No   No  No   No
  12    No    Yes  True  No   No   No  No   No
  21    No    Yes  True  No   No   No  No   No
  Trk1  No    Yes  True  No   No   No  No   No
```

Viewing the global RPVST+ configuration per port per VLAN

Syntax:

```
show spanning-tree<ethernet>port-list vlan vlan-id
```

Lists the spanning tree port parameter settings per port per VLAN.

Viewing the global RPVST+ configuration per port per VLAN

```
switch# show spanning-tree 9 config vlan 10

 Spanning Tree Information

  STP Enabled          [No] : Yes
  Mode                      : RPVST
  Extended System ID        : Enabled
  Ignore PVID Inconsistency : Disabled
  Switch MAC Address        : 002347-587b80


  RPVST Enabled             : Enabled
  VLAN ID                   : 10
  Switch Priority           : 32768
  Forward Delay             : 15
  Hello Time                : 2
  Max Age                   : 20
  Admin Root Bridge         : Not Configured

       Path      Port     Admin Auto Admin Root Loop TCN BPDU BPDU
 Port  Cost      Priority Edge  Edge PtP   Grd  Grd  Grd Flt  Guard
 ----- ---------- -------- ----- ---- ----- ---- ---- --- ---- -----
 9     20000     128       No    Yes  True  No   No   No  No   No
```

Viewing the global RPVST+ configuration per VLAN

Syntax:

```
show spanning-tree config vlan vlan-id
```

Lists the spanning tree port parameter settings for only the specified VLAN.

Viewing the global RPVST+ configuration per VLAN

```
switch(config)# show spanning-tree config vlan 20

 Spanning Tree Information

  STP Enabled          [No] : Yes
  Mode                      : RPVST
  Extended System ID        : Enabled
  Ignore PVID Inconsistency : Disabled
  Switch MAC Address        : 002347-587b80


  RPVST Enabled             : Enabled
  VLAN ID                   : 20
  Switch Priority           : 32768
  Forward Delay             : 15
  Hello Time                : 2
  Max Age                   : 20
  Admin Root Bridge         : Not Configured

                   Path      Port
  Port  Type       Cost      Priority
  ----- ---------- --------- ---------
  9     100/1000T  20000     128
  20    100/1000T  200000    128
  21    100/1000T  20000     128
```

Viewing BPDU status and related information

Syntax:

```
show spanning-tree bpdu-protection port-list
```

Displays the BPDU protection state and errant BPDU count for ports in the port list.

Viewing BPDU status in show spanning tree output

```
switch# show spanning-tree 22

 Spanning Tree Information

  STP Enabled    [No] : Yes
  Mode                : RPVST
  RPVST Enabled VLANs : 10,20


  Switch MAC Address   : 0024a8-d13a40

  Port                   : 22
  Status                 : Up           Port Type           : 100/1000T
  BPDU Protection        : No           BPDU Filtering      : No
  Root Guard             : No           TCN Guard           : No
  Loop Guard             : No           Admin Edge Port     : No
  Admin PointToPoint MAC : Yes

  VLAN   Port       Port      Port       Designated    Hello Oper  Oper
  ID     Path-Cost  Priority  State      Bridge        Time  Edge  PtP
  ------ ---------- --------- ---------- ------------- ----- ----- ------
  20     20000      128       Forwarding 0024a8-d13a40 2     No    Yes
```

Viewing BPDU protection status on specific ports

```
switch# show spanning-tree bpdu-protection 11-12,21-24

 Status and Counters - STP BPDU Protection Information

 BPDU Protection Timeout (sec) : 60
 BPDU Protected Ports : 23-24


  Port   Type      Protection State             Errant BPDUs
  ------ --------- ---------- ----------------- ------------
  11     100/1000T No                           0
  12     100/1000T No                           0
  21     100/1000T No                           0
  22     100/1000T No                           0
  23     100/1000T Yes                          0
  24     100/1000T Yes                          0
```

Viewing RPVST+ VLAN and vPort system limits

Each switch model supports a maximum number of active virtual ports (vPorts). New port VLAN memberships cannot be created once the vPort limit has been reached. Also, there is a maximum recommended number of active vPorts for each fixed-port switch or each module in a chassis switch. Exceeding the maximum recommended number of vPorts can result in dropped BPDUs and potential network loops. This command displays the current vPort status and maximum recommended vPort total per-switch or, for modular switches, per-module.

Syntax:

```
show spanning-tree system-limits rapid-pvst
```

Displays the RPVST+ VLAN and virtual port (vPort) status on the switch.

Viewing RPVST+ VLAN and vPort system limits

```
switch(config)# show spanning-tree system-limits rapid-pvst

 Spanning Tree Information


  STP Enabled               : Yes
  Mode                      : RPVST
  RPVST Enabled VLANs       : 20


  Switch MAC Address                    : 002347-c651c0
  Count of RPVST Enabled VLANs          : 1
  Maximum Allowed RPVST Enabled VLANs   : 400
  Count Of Total Virtual Ports          : 24
  Maximum Allowed Virtual Ports         : 424

                       Current          Operational      Recommended Maximum
  Ports                Virtual Ports    Virtual Ports    Virtual Ports
  -------------------- ---------------- ---------------- --------------------
  Ports 1-24           24               2                200
```

Table 28: *Virtual Port Data Fields*

vPort data field	Description
`Count of Total Virtual Ports`	The count of active vPorts (ports per VLAN) plus the count of non-active vPorts (all ports that belong to trunks).
`Maximum Allowed Virtual Ports`	The total of the system-created vPort instances plus the maximum user-assignable vPort instances. Each port on the switch belongs to at least one VLAN (VLAN-1 by default), which is a system-created vPort instance. The user-assigned VPORT instances are the system-assigned vPort instances. The `show spanning-tree system-limits rapid-pvst` command combines the system-created vPort instances and the user-assigned maximum vPort instances when calculating the maximum allowed virtual ports. **NOTE:** Each user-configured trunk on the switch increments this value by 1.
`Current Virtual Ports`	The number of ports that are members of each VLAN on a per-module basis (or a per-group of ports basis).
`Operational Virtual Ports`	The number of ports belonging to each PVST-enabled VLAN on a per-module basis (or a per-group of ports basis). This value should not exceed the recommended maximum vPort limit.
`Recommended Maximum Virtual Ports`	The maximum recommended number of vPort instances that should be allowed on the switch. Exceeding this limit can potentially result in received BPDUs being dropped.

Configuring vPorts

Virtual ports on a switch are calculated as ports per-VLAN. Also, a trunk membership on one or more VLANs counts as one vPort per-VLAN, regardless of how many physical ports belong to the trunk. For example, the following configuration on a modular chassis results in 26 vPorts.

```
trunk 1,2 trk1
vlan 1
   name "DEFAULT_VLAN"
   untagged 3-24
   no untagged trk1
   exit
vlan 20
   ip address 10.243.230.75 255.255.255.248
   name "VLAN20"
   tagged trk1
   exit
vlan 30
   ip address 10.243.230.83 255.255.255.248
   name "VLAN30"
   tagged 13,14,trk1
   exit
```

	Module “A”	Module “B”	Module “C”	Total vPorts on the Switch
VLAN 1	22 (A3 - A24)	23 (B2 - B24	24 (C1 - C24)	
VLAN 20	1 (trk1: A1 - A2)	1 (trk1: B1)[1]	0	
VLAN 30	2 (A13 - A14) 1 (trk1: A1 - A2)[1]	2 (B13 - B14) 1 (trk1: B1)[1]	0	
vPorts per-module	26	27	24	77

[1] A trunk in a given VLAN counts as one vPort for each module on which it occurs.

Calculating non-active vPorts

Every port that is part of a manually configured trunk is counted as a non-active (reserved) vPort. For example, the ports in the following configuration are all non-active vPorts:

```
trunk 1, 2 trk1
trunk 3-5 trk2 lacp
trunk 17-20 trk3 dt-lacp
```

Calculating per-module vPorts on chassis switches

The switch-wide active vPort count, there is a vPort count per port module determined by the number of ports per line card that are members of each VLAN. Also, on modular switches, if a VLAN includes a trunk configured with ports on more than one module, then one vPort is counted for each module on which the trunk exists (regardless of how many ports are included in the trunk.) For example, in the following configuration, VLANs 1, 20, and 30 have a total of 74 vPorts.

```
trunk A1,A2,B1 trk1
vlan 1
    name “DEFAULT_VLAN”
    untagged A3-A24, B2-B24
    no untagged trk1
    exit
vlan 20
    ip address 10.243.230.75 255.255.255.248
    name “VLAN20”
    tagged A3-A12, B2-B12, trk1
    exit
vlan 30
    ip address 10.243.230.83 255.255/255/248
    name “VLAN30”
    tagged A13, A14, B13, B14, trk1
    exit
```

	Module “A”	Module “B”	All Modules
VLAN 1	22	23	4
VLAN 20	10 + 1	11 + 1	23
VLAN 30	2 + 1	2 + 1	6
Total vPorts	36	38	74

Troubleshooting an RPVST+ configuration

Cause

This section describes the show spanning tree commands you can use to monitor, troubleshoot, and debug the operation of a per-VLAN spanning tree configuration in your network.

NOTE: The `show spanning-tree`commands described in this section allow you to troubleshoot RPVST+ activity in your network by focusing on increasingly specific levels of operation. For example, you can display debug information for:

- All VLANs
- All ports of one VLAN
- A specific port or several ports used in one VLAN

Viewing the change history of root bridges

Syntax:

```
show spanning-tree root-history vlan vlan-id
```

Displays the last 10 root bridge changes on a specified VLAN configured with RPVST+. Included are the timestamp and Root Bridge ID recorded at each root bridge change.

Use the `show spanning-tree root-history` command to view the number and dates of changes in the assignment of a root bridge. Possible intrusion into your VLAN network may occur if an unauthorized external device gains access to a spanning tree by posing as the root device in a topology. To prevent a port connected to the device from being selected as the root port in a topology, use the `spanning-tree root-guard` command.

Viewing the change history of root bridges

```
switch# show spanning-tree root-history vlan 20

 Status and Counters - RPVST Root Changes History

  VLAN ID                 : 20
  Root Changes Counter    : 53
  Current Root Bridge ID : 32768:0024a8-d13a40

 Root Bridge ID       Date        Time
```

```
-------------------- ---------- --------
32768:0024a8-d13a40 05/04/2012 21:54:11
    0:001185-c6e500 05/04/2012 21:54:07
32768:0024a8-d13a40 05/04/2012 16:41:11
    0:001185-c6e500 05/04/2012 16:41:11
```

Enabling traps and Viewing trap configuration

Syntax:

`no spanning-tree trap`

`[[errant-bpdu | loop-guard | new-root] | [topology-change [vlan<vid-list | instance [instance-ID] {cstt}|root-guard]]`

Enables or disables SNMP traps for errant-BPDU, loop guard, new root, and root guard event notifications.

`errant-bpdu`

Enables SNMP notification when an errant BPDU is received. Designed for use with BPDU filtering.

`loop-guard`

Enables SNMP notification when a loop guard inconsistency is detected. Designed for use with the Loop Guard option.

`new-root`

Enables SNMP notification when a new root is elected on any VLAN configured for RPVST+ on the switch.

`root-guard`

Enables SNMP notifications when a root-guard inconsistency is detected.

`topology-change`

Enables notifications sent when a topology change occurs.

`topology-change-history`

Shows the spanning tree topology history changes.

Default for all of the above options: Disabled

The `no` form of the command disables traps on the switch.

Syntax:

`show spanning-tree traps`

Displays the current spanning tree trap configuration on the switch.

Viewing spanning tree traps in the default configuration

```
switch(config)# show spanning-tree traps

 Status and Counters - STP Traps Information

 Trap Name                | Status
 ---------------------- + --------
 errant-bpdu              | Disabled
 new-root                 | Disabled
 root-guard               | Disabled
 loop-guard               | Disabled
```

Viewing debug counters for all VLAN instances

Syntax:

```
show spanning-tree debug-counters
```

Displays the aggregate values of all RPVST+ debug counters that are maintained on a switch. These aggregate values are a summary of the information collected from all ports and from all spanning tree instances for all switch ports. Use the displayed diagnostic information to globally monitor RPVST+ operation on a per-switch basis.

Viewing debug counters for all VLANs

```
switch# show spanning-tree debug-counters

 Status and Counters - RPVST Debug Counters Information

                               Aggregated Value
  Counter Name                 Collected from VLANs
  ------------------------------ -------------------------
  Invalid BPDUs                  0
  Errant BPDUs                   0
  Looped-back BPDUs              0
  Starved BPDUs                  18
  Exceeded Max Age BPDUs         3
  Topology Changes Detected      9
  Topology Changes Tx            9
  Topology Changes Rx            4
  Topology Change ACKs Tx        0
  Topology Change ACKs Rx        6
  TCN BPDUs Tx                   4
  TCN BPDUs Rx                   0
  CFG BPDUs Tx                   0
  CFG BPDUs Rx                   0
  RST BPDUs Tx                   0
  RST BPDUs Rx                   0
  RPVST BPDUs Tx                 1881
  RPVST BPDUs Rx                 2617
```

Viewing debug counters per-VLAN

Syntax:

```
show spanning-tree debug vlan vlan-id
```

Displays the aggregate values of all RPVST+ debug counters maintained on a switch for a specified VLAN.

Viewing debug counters for a specific VLAN

```
switch(config)# show spanning-tree debug vlan 20

 Status and Counters - RPVST Debug Counters Information

  VLAN ID : 20

                                 Aggregated Value
  Counter Name                   Collected from Ports
  ------------------------------ --------------------
  Invalid BPDUs                  5
  Errant BPDUs                   10
  Looped-back BPDUs              0
  Starved BPDUs                  9
  Exceeded Max Age BPDUs         2
  Topology Changes Detected      9
  Topology Changes Tx            4
  Topology Changes Rx            181
  Topology Change ACKs Tx        0
  Topology Change ACKs Rx        0
  TCN BPDUs Tx                   0
  TCN BPDUs Rx                   0
  CFG BPDUs Tx                   0
  CFG BPDUs Rx                   0
  RST BPDUs Tx                   0
  RST BPDUs Rx                   0
  RPVST BPDUs Tx                 1531
  RPVST BPDUs Rx                 1428
```

Viewing debug counters per-port per-VLAN

Syntax:

```
show spanning-tree debug ports port-list vlan vlan-id
```

Displays the aggregate values of all RPVST+ debug counters maintained on one or more ports used by a specified VLAN.

Viewing debug counters for a specific port on a VLAN

```
Switch_A(config)# show spanning-tree debug ports 9 vlan 20

 Status and Counters - RPVST Debug Counters Information

  VLAN ID : 20
  Port : 9

  Counter Name                   Value           Last Updated
  ------------------------------ --------------- --------------------
  Invalid BPDUs                  0               04/16/2012 22:27:15
  Errant BPDUs                   0               04/16/2012 22:27:15
  Looped-back BPDUs              0               04/16/2012 22:27:15
  Starved BPDUs                  5               05/01/2012 21:48:11
  Exceeded Max Age BPDUs         0               04/16/2012 22:27:15
  Topology Changes Detected      9               05/04/2012 21:54:05
  Topology Changes Tx            5               05/05/2012 22:04:49
  Topology Changes Rx            2               05/07/2012 18:08:34
  Topology Change ACKs Tx        0               04/16/2012 22:27:15
  Topology Change ACKs Rx        0               04/16/2012 22:27:15
  TCN BPDUs Tx                   0               04/16/2012 22:27:15
  TCN BPDUs Rx                   0               04/16/2012 22:27:15
  CFG BPDUs Tx                   0               04/16/2012 22:27:15
  CFG BPDUs Rx                   0               04/16/2012 22:27:15
  RST BPDUs Tx                   0               04/16/2012 22:27:15
  RST BPDUs Rx                   0               04/16/2012 22:27:15
  RPVST BPDUs Tx                 7812            05/05/2012 22:04:49
  RPVST BPDUs Rx                 1065            05/08/2012 19:43:11
```

Field descriptions for RPVST+ debug command output

Field	Shows the number of —
`Invalid BPDUs`	Received BPDUs that failed standard RPVST+ (802.1Q-REV/D5.0 14.4) validation checks and were dropped. This counter is maintained on a per-port per-VLAN basis.
`Errant BPDUs`	Received BPDUs that were dropped on a port that is configured to not expect BPDU packets. This counter is maintained on a per-port basis and is incremented each time a BPDU is received on a port configured with the BPDU filter to ignore incoming BPDU packets (`spanning-tree bpdu-filter` command) or the BPDU protection feature to disable the port when BPDU packets are received (`spanning-tree bpdu-protection` command).
`Looped-back BPDUs`	Times that a port has received self-sent BPDU packets as the result of an external loop condition in which the BPDUs were looped back to the originating transmission port. The received BPDU is still processed by RPVST+ and the port changes to a blocked state. This counter is maintained on a per-port per-VLAN basis.
`Starved BPDUs`	Times that no BPDUs are received within the scheduled interval (three times the Hello Time value configured with the `spanning-tree vlan hello-time` command) from a VLAN-designated peer port on the VLAN root, alternate, or backup port. As a result, the "starved" port triggers a spanning tree topology regeneration. This counter is maintained on a per-port per-VLAN basis.
`Exceeded Max Age BPDUs`	Times that a BPDU packet is received from a bridge with a Message Age value greater than the configured value of the Max Age parameter `(spanning-tree maximum age` command). This may occur if the receiving bridge is located too far from the root bridge (beyond the configured size of the spanning tree domain on the root bridge) or if a BPDU packet with invalid root information is continuously circulating between bridges in a spanning tree domain and needs to be aged out.
`Topology Changes Detected`	Times that a Topology Change event is detected by the port on a given VLAN and the port triggers a topology change propagation throughout the network. A Topology Change event occurs when a non-edge port enters forwarding state. This counter is maintained on a per-VLAN per-port basis.
`Topology Changes Tx`	Times that Topology Change information is propagated (sent out) through the port to the rest of the network. For a VLAN port running PVST (non-rapid), the counter is the number of times that a CFG or RST BPDU with the TC flag set is transmitted out of the port. This counter is maintained on a per-VLAN per-port basis.
`Topology Changes Rx`	Times that Topology Change information is received from the peer port. For a VLAN port running PVST (non-rapid), the counter is the number of times that a CFG or RST BPDU with the TC flag set is received. This counter is maintained on a per-port per-VLAN basis.
`Topology Change ACKs Tx`	Times that the Topology Change acknowledgement is transmitted through the port (number of CFG or RST BPDUs transmitted with the Topology Change Acknowledge flag set). This counter is maintained on a per-port per-VLAN basis.

Table Continued

Field	Shows the number of —
`Topology Change ACKs Rx`	Times that the Topology Change acknowledgement is received on the port (number of CFG or RST BPDUs received with the Topology Change Acknowledge flag set). This counter is maintained on a per-VLAN basis.
`TCN BPDUs Tx`	Topology Change Notification BPDUs that are transmitted through the port. This counter is maintained on a per-port basis.
`TCN BPDUs Rx`	Topology Change Notification BPDUs that are received on the port. This counter is maintained on a per-port per-VLAN basis.
`CFG BPDUs Tx`	802.1D configuration BPDUs that are transmitted through the port. This counter is maintained on a per-port per-VLAN basis.
`CFG BPDUs Rx`	802.1D configuration BPDUs that are received on the port. This counter maintained on a per-port per-VLAN basis.
`RST BPDUs Tx`	802.1w RST BPDUs that are transmitted through the port. This counter is maintained on a per-port per-VLAN basis.
`RST BPDUs Rx`	802.1w RST BPDUs that are received on the port. This counter is maintained on a per-port per-VLAN basis.

RPVST+ event log messages

Event	Log message
STP enabled/disabled on a VLAN	Spanning tree Protocol enabled/disabled on vlan *vlan-id*
Switch does not receive BPDUs from peer on a particular VLAN and port	VLAN *vlan-id* starved for a BPDU on port *port number* from *bridge name*
Switch received BPDU with inconsistent VLAN	Blocking port-name on vlan *vlan-id*.
Inconsistency is restored	Unblocking port-name on vlan *vlan-id* Port consistency restored.
Root port is changed on a VLAN	VLAN *vlan-id*root changed from *bridgepriorty:mac* to *bridge priority:mac*
Switch received a BPDU with invalid TLV	Received SSTP BPDU with bad TLV on *port-number vlan-id*
The number of `vlan-port` instances exceeds the recommended limit	The number of vlan-port instances exceeded the recommended limit of *num*

Table Continued

Event	Log message
RADIUS subsystem tries to dynamically change port VLAN assignments when mode is RPVST	`RADIUS unable to assign port to VLAN` *`vlan-id`* `because spanning-tree is running in RPVST+ mode`
LLDP subsystem tries to dynamically change port VLAN assignments when mode is RPVST	`LLDP unable to assign port` *`port-number`* `to VLAN` *`vlan-id`* `because spanning-tree is running in RPVST+ mode`
VPORT counts exceed 200	`The number of vPorts on slot` *`slot-number`* `exceeds the recommended limit of` *`vport-count`*`. PVST BPDUs may be dropped.`

Using RPVST+ debug

While the Event Log records switch-level progress, status, and warning messages on the switch, the Debug/ System Logging (Syslog) feature provides a way to record Event Log and debug messages on a remote device. For example, you can send messages about routing misconfigurations and other network protocol details to an external device, and later use them to debug network-level problems. The Debug/Syslog feature allows you to specify the types of Event Log and debug messages that you want to send to an external device. The two commands described next affect debug operation for RPVST+. For further information on debug operation, see the *ArubaOS-Switch Management and Configuration Guide* for your switch.

Syntax:

```
spanning-tree clear-debug-counters [ports port-list] [vlan vid-list]
```

Clears all spanning tree debug counters unless specific ports and VLANs are specified.

`ports` ***`port-list`***

Clears spanning tree debug counters on the specified ports.

`[ethernet]`***`port-list`***

Clears spanning tree debug counters on an entered list of posts or `all` for the `ports` command parameter.

`vlan`

Clears spanning tree debug counters for the VLAN.

`vlan` ***`vlan-id-list`***

One or more identifiers for the `VLAN` command parameter.

Using the `vlan` and `ports` options together clears the spanning tree debug counters on the specified ports for the specified VLANs. Counters maintained on the same ports for other VLANs are not cleared.

Syntax:

```
debug rpvst [event [filter vlan vid-list]]
```

```
no debug rpvst [event [filter vlan vid-list]]
```

```
debug rpvst [packet [filter port port-list [vlan vid-list]]]
```

```
no debug rpvst [packet [filter port port-list [vlan vid-list]]]
```

Displays RPVST+ debug messages on the destination device specified with the `debug destination logging | session | buffer` command.

event

Displays RPVST+ Event Log messages.

filter vlan *vid-list*

Limits log messages to those generated on the specified VLANs.

packet

Displays RPVST+ packets sent and received.

filter port *port-list* vlan *vid-list*

Limits packets displayed to those generated on the specified ports. If the `vlan` option is used, then packets displayed are further limited to the ports on the specified VLANs.

The `no` form of the command disables display of RPVST+ debug messages on the destination device.

Introduction to switch management

Stacking enables you to use a single IP address and standard network cabling to manage a group of up to 16 total switches in the same IP subnet (broadcast domain). Using stacking, you can:

- Reduce the number of IP addresses needed in your network.
- Simplify management of small workgroups or wiring closets while scaling your network to handle increased bandwidth demand.
- Eliminate any specialized cables for stacking connectivity and remove the distance barriers that typically limit your topology options when using other stacking technologies.
- Add switches to your network without having to first perform IP addressing tasks.

NOTE:

- When an older model of switch is running as the commander of a stack, the web GUI cannot be used to monitor or manage a stack member that is of a newer switch generation.
- In the default configuration, stacking in the "candidate" state is enabled.
- For additional rules and restrictions, see **Operating rules for stacking**.

Configuring stack management

This process assumes that:

- All switches to include in a stack are connected to the same subnet (broadcast domain).
- If VLANs are enabled on the switches to include in the stack, then the ports linking the stacked switches must be on the primary VLAN in each switch. If the primary VLAN is tagged, then each switch in the stack must use the same VLAN ID (VID) for the primary VLAN.

Options for configuring a commander and candidates

Depending on how Commander and Candidate switches are configured, Candidates can join a stack either automatically or by a Commander manually adding (pulling) them into the stack. In the default configuration, a Candidate joins only when manually pulled by a Commander, but you can reconfigure a Commander to automatically pull in Candidates that are in the default stacking configuration. Also a Candidate switch can be re-configured to either "push" itself into a particular Commander's stack, convert the Candidate to a Commander (for a stack that does not already have a Commander), or to operate as a standalone switch without stacking. The following table shows your control options for adding Members to a stack.

Table 29: *Stacking configuration guidelines*

Join Method 1	Commander (IP Addressing Required)Auto Grab	Candidate (IP Addressing Optional)Auto Join	Passwords
Automatically add Candidate to Stack Causes the first 15 eligible, discovered switches in the subnet to automatically join a stack.	Yes	Yes (default)	No (default)
Manually add Candidate to StackPrevent automatic joining of switches you don't want in the stack	No (default)	Yes (default)	Optional[1]
	Yes	No	Optional[1]
	Yes	Yes (default) or No	Configured
Prevent a switch from being a Candidate	N/A	Disabled	Optional

[1] The Commander's Manager and Operator passwords propagate to the candidate when it joins the stack.

The easiest way to automatically create a stack is to:

Procedure

1. Configure a switch as a Commander.
2. Configure IP addressing and a stack name on the Commander.
3. Set the Commander's `Auto Grab` parameter to `Yes`.
4. Connect Candidate switches (in their factory default configuration) to the network.

This approach automatically creates a stack of up to 16 switches (including the Commander). However this replaces manual control with an automatic process that may bring switches into the stack that you did not intend to include. With the Commander's `Auto Grab` parameter set to `Yes`, any switch conforming to all four of the following factors automatically becomes a stack Member:

- Default stacking configuration (Stack State set to Candidate, and Auto Join set to Yes)
- Same subnet (broadcast domain) and default VLAN as the Commander (If VLANs are used in the stack environment, see "Stacking Operation with Multiple VLANs Configured” on page **Stacking operation with multiple VLANs configured** on page 266.)
- No Manager password
- 14 or fewer stack members at the moment

Creating a stack (Overview)

Procedure

1. Determine the naming conventions for the stack.

   A stack name is necessary. To help distinguish one switch from another in the stack, configure a unique system name for each switch. Otherwise, the system name for a switch appearing in the Stacking Status screen appears as the stack name plus an automatically assigned switch number. For example:

   Figure 40: *Using the system name to help identify individual switches*

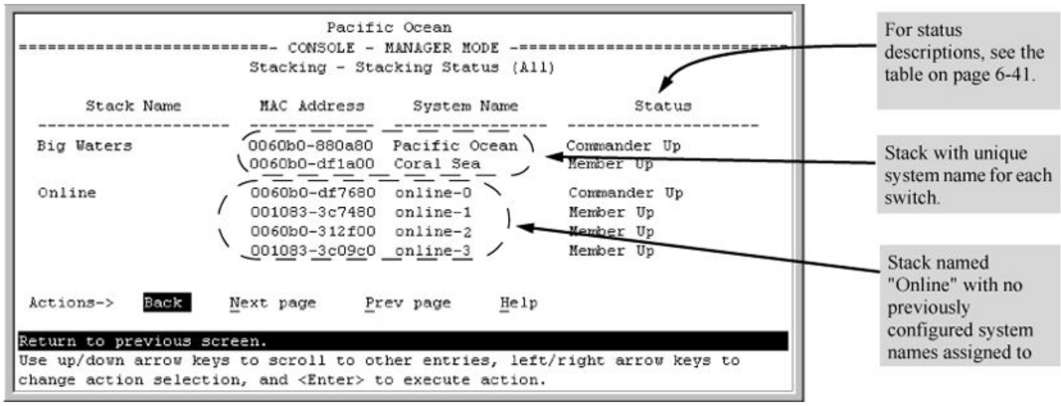

```
                              Pacific Ocean
=========================- CONSOLE - MANAGER MODE -===========================
                      Stacking - Stacking Status (All)

      Stack Name          MAC Address     System Name           Status
  --------------------  -------------  --------------   --------------------
  Big Waters            0060b0-880a80  Pacific Ocean    Commander Up
                        0060b0-df1a00  Coral Sea        Member Up
  Online                0060b0-df7680  online-0         Commander Up
                        001083-3c7480  online-1         Member Up
                        0060b0-312f00  online-2         Member Up
                        001083-3c09c0  online-3         Member Up

 Actions->   Back     Next page     Prev page     Help

Return to previous screen.
Use up/down arrow keys to scroll to other entries, left/right arrow keys to
change action selection, and <Enter> to execute action.
```

For status descriptions, see the table on page 6-41.

Stack with unique system name for each switch.

Stack named "Online" with no previously configured system names assigned to

2. Configure the Commander switch.

   Doing this first helps establish consistency in your stack configuration, which can help prevent startup problems.

   - A stack requires one Commander switch. If you plan to implement more than one stack in a subnet, to avoid unintentionally adding a Candidate to the wrong stack, manually control the joining process by leaving the Commander's `Auto Grab` parameter set to `No`.
   - The Commander assigns its Manager and Operator passwords to any Candidate switch that joins the stack.
   - The Commander's SNMP community names apply to members.

3. For automatically or manually pulling Candidate switches into a stack, leave such switches in their default stacking configuration; to access Candidate switches through your network before they join the stack, assign IP addresses to these devices. Otherwise, IP addressing is optional for Candidates and Members.

   NOTE: Once a Candidate becomes a member, you can access it through the Commander to assign IP addressing or make other configuration changes.

4. Make a record of any Manager passwords assigned to the switches intended for your stack, that are not currently members. Using these passwords enables the protected switches to join the stack.
5. If using VLANs in the stacking environment, use the default VLAN for stacking links. See **Stacking operation with multiple VLANs configured** on page 266.
6. Ensure that all switches intended for the stack are connected to the same subnet. As soon as you connect the Commander, it begins discovering available Candidates in the subnet.

   a. If the Commander is configured to automatically add Members (`Auto Grab=Yes`), the first fifteen discovered Candidates meeting both the following criteria will automatically join the stack:

  - `Auto Join` parameter set to `Yes` (the default)
  - Manager password not configured

b. If you configured the Commander to manually add Members (`Auto Grab` set to `No` - the default), begin the process of selecting and adding the desired Candidates.

7. Ensure that all switches intended for the stack have joined.

8. If you need to do specific configuration or monitoring tasks on a Member, use the console interface on the Commander to access the Member.

Viewing stack status (CLI)

Syntax:

```
show stack [candidates | view | all]
```

Lists the stack status for an individual switch or other switches discovered in the same subnet.

Viewing the status of an individual switch

Syntax:

```
show stack
```

Lists the stacking configuration for an individual switch.

Displaying `show stack` output

```
switch(config)# show stack

Stacking - Stacking Status (This Switch)
 Stack State            : Commander
 Transmission Interval : 60
 Stack Name             : Big_Waters    Number of members         : 14
 Auto Grab              : Yes           Members unreachable       : 0

 SN MAC Address    System Name    Device Type           Status
 -- ------------- ------------- --------------------- -----------------
 0 1cc1de-cfbc80  Big_Waters-0  switch                Commander Up
 1 000883-08f980  Big_Waters-1  switch                Member Up
```

Viewing the status of candidates the Commander has detected (CLI)

Syntax:

```
show stack candidates
```

Using the `show stack candidates` command to list candidates

```
switch(config)# show stack candidates
 Stack Candidates
```

```
  Candidate MAC System Name            Device Type
  ------------- ---------------------- ---------------------
  0060b0-889e00 DEFAULT_CONFIG         switch
```

Viewing the status of all stack-enabled switches discovered in the IP subnet (CLI)

Syntax:

```
show stack all
```

Lists all the stack-configured switches discovered in the IP subnet.

Using the `show stack all` command to list discovered switches in the IP subnet

In this example, because the switch on which the `show stack all` command was executed is a Candidate, it is included in the Others category.

```
switch(config)# show stack all

Stacking - Stacking Status (All)

 Stack Name      MAC Address   System Name               Status
 --------------- ------------- ------------------------- --------------------
 Big_Waters      1cc1de-cfbc80 Big_Waters-0              Commander Up
                 000883-08f980 Big_Waters-1              Member Up
 Others:         0060b0-889e00 DEFAULT_CONFIG            Candidate
```

Viewing the status of the Commander and current members of the Commander’s stack (CLI)

Syntax:

```
show stack view
```

Lists all switches in the stack of the selected switch.

Using the `show stack view` command to list the stack assigned to the selected commander

```
switch(config)# show stack view
Stack Members

 SN MAC Address   System Name   Device Type          Status
 -- ------------- ------------- -------------------- ------------
 0 1cc1de-cfbc80  Big_Waters-0  switch               Commander Up
 1 000883-08f980  Big_Waters-1  switch               Member Up
```

Configuring a Commander switch (CLI)

Any stacking-enabled switch can become a Commander as long as the intended stack name does not already exist on the broadcast domain. This is because creating a Commander automatically creates a corresponding stack.

Before you begin configuring stacking parameters:

Procedure

1. Configure IP addressing on the switch intended for stack commander and, if not already configured, on the primary VLAN. (For more on configuring IP addressing, see the *ArubaOS-Switch Management and Configuration Guide* for your switch.)

   NOTE: The primary VLAN must have an IP address in order for stacking to operate properly. For more on the primary VLAN, see *Aruba-OS Switch Management and Configuration Guide* for your switch.

2. Configure a Manager password on the switch intended for Commander; (the Commander's Manager password controls access to stack Members). For more on passwords, see the local manager and operator password information in the *ArubaOS-Switch Access Security Guide* for your switch.

Making a switch a Commander (CLI)

Syntax:

```
stack commander <name-str>
```

Assigns a stack name to a switch, makes it a Commander and automatically creates a stack.

Creating a Commander switch

To create a Commander switch with a stack name of `Big_Waters`:

```
switch(config)# stack commander Big_Waters
```

NOTE: If stacking was previously disabled on the switch, this command also enables stacking.

As the following `show stack` display shows, the Commander switch is now ready to add members to the stack.

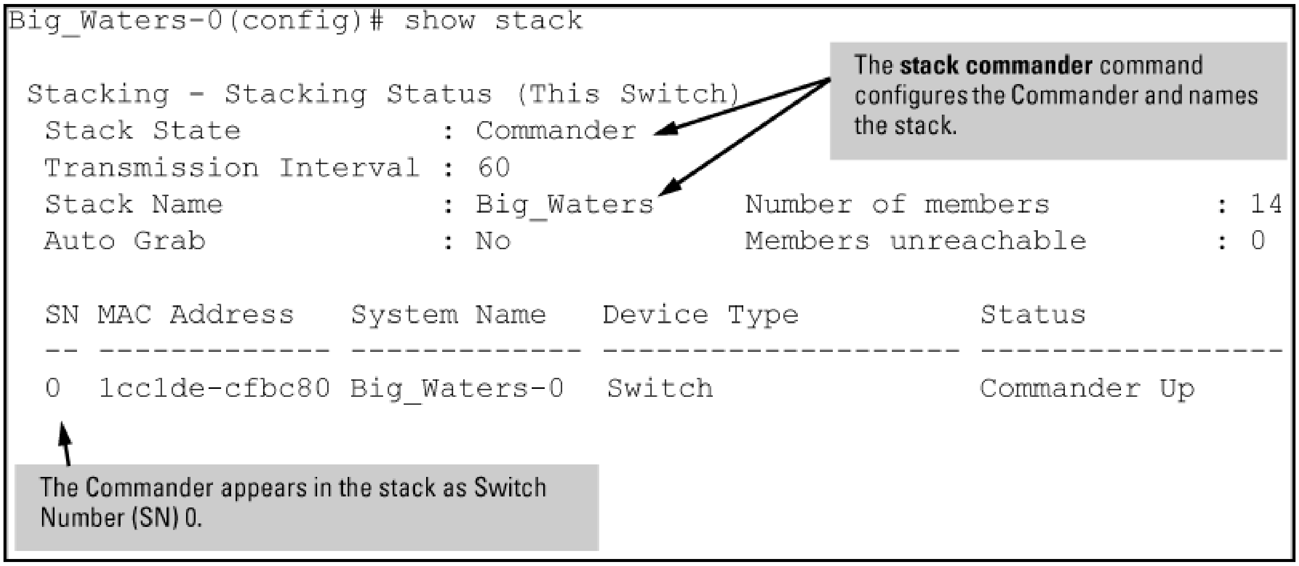

```
Big_Waters-0(config)# show stack

 Stacking - Stacking Status (This Switch)
  Stack State            : Commander
  Transmission Interval  : 60
  Stack Name             : Big_Waters    Number of members       : 14
  Auto Grab              : No            Members unreachable     : 0

  SN MAC Address   System Name   Device Type          Status
  -- ------------- ------------- -------------------- ----------------
  0  1cc1de-cfbc80 Big_Waters-0  Switch               Commander Up
```

The **stack commander** command configures the Commander and names the stack.

The Commander appears in the stack as Switch Number (SN) 0.

Using a Member's CLI to make the Member Commander of a new stack

This procedure requires that you first remove the Member from its current stack, then create the new stack.

NOTE: If you do not know the MAC address for the Commander of the current stack, use `show stack` to list it.

Syntax:

```
no stack join <MAC address>
```

```
stack commander <stack name>
```

Suppose a switch named `Bering Sea` is a Member of a stack named `Big_Waters`. To use the CLI to convert it from stack Member to Commander of a new stack named "Lakes", use the above commands as shown in the following example.

Using a member's CLI to convert the member to the commander of a new stack

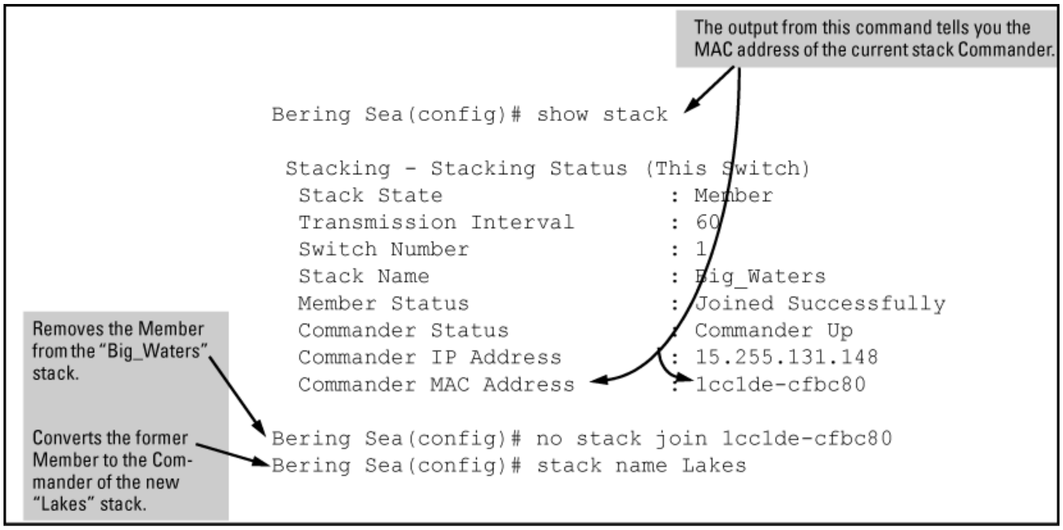

The output from this command tells you the MAC address of the current stack Commander.

```
Bering Sea(config)# show stack

 Stacking - Stacking Status (This Switch)
  Stack State                 : Member
  Transmission Interval       : 60
  Switch Number               : 1
  Stack Name                  : Big_Waters
  Member Status               : Joined Successfully
  Commander Status            : Commander Up
  Commander IP Address        : 15.255.131.148
  Commander MAC Address       : 1cc1de-cfbc80

Bering Sea(config)# no stack join 1cc1de-cfbc80
Bering Sea(config)# stack name Lakes
```

Removes the Member from the "Big_Waters" stack.

Converts the former Member to the Commander of the new "Lakes" stack.

Adding to a stack, or moving switches between stacks (CLI)

You can add switches to a stack by adding discovered Candidates or by moving switches from other stacks that may exist in the same subnet.

NOTE: You cannot add a Candidate that the Commander has not discovered.

In its default configuration, the Commander's `Auto-Grab` parameter is set to `No` to give you manual control over which switches join the stack and when they join. This prevents the Commander from automatically trying to add every Candidate it finds that has `Auto Join` set to `Yes` (the default for the Candidate).

To allow eligible candidates to automatically join the stack when the Commander discovers it, configure `Auto Grab` in the Commander to `Yes`. When you do so, any Candidate discovered with `Auto Join` set to `Yes` (the default) and no Manager password, will join the stack, up to the limit of 15 Members.

To manually add a candidate, use:

- A switch number (SN) to assign to the new member. Member SNs range from `show stack view`. You can use any SN not included in the listing. (SNs are viewable only on a Commander switch.)
- The MAC address of the discovered Candidate you are adding to the stack; to see the MAC address, use `show stack view`.

Finding switch numbers (SNs)

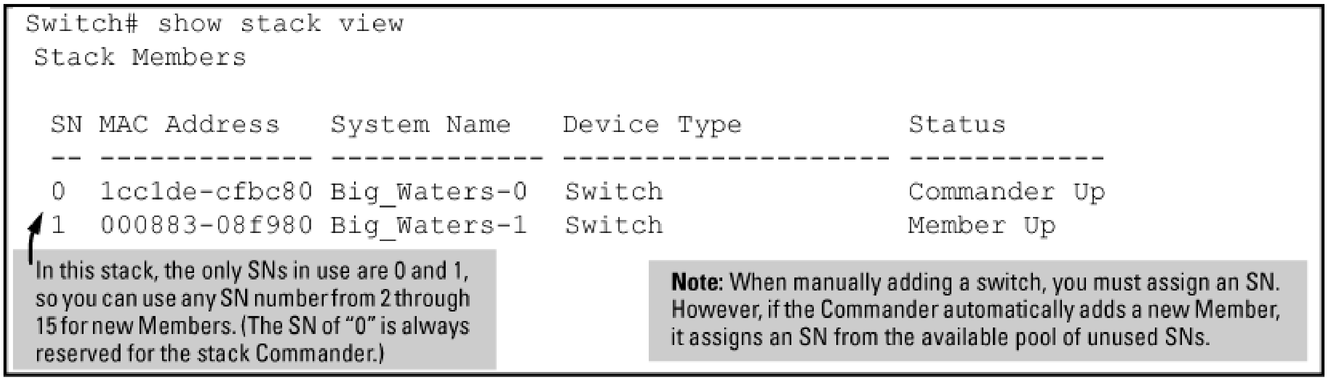

```
Switch# show stack view
 Stack Members

   SN MAC Address   System Name   Device Type          Status
   -- ------------- ------------- -------------------- ------------
   0  1cc1de-cfbc80 Big_Waters-0  Switch               Commander Up
   1  000883-08f980 Big_Waters-1  Switch               Member Up
```

In this stack, the only SNs in use are 0 and 1, so you can use any SN number from 2 through 15 for new Members. (The SN of "0" is always reserved for the stack Commander.)

Note: When manually adding a switch, you must assign an SN. However, if the Commander automatically adds a new Member, it assigns an SN from the available pool of unused SNs.

To display all discovered Candidates with their MAC addresses, use the `show stack candidates` command from the Commander CLI. For example, to list discovered candidates for the above Commander:

Finding MAC addresses

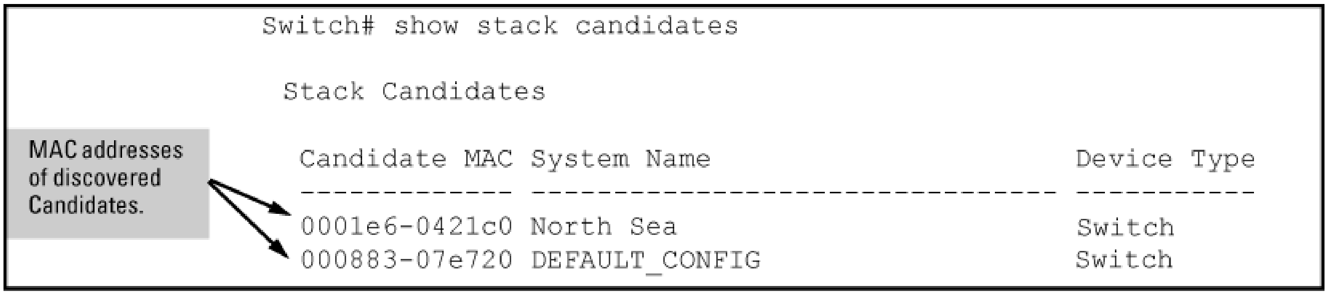

```
Switch# show stack candidates

 Stack Candidates

  Candidate MAC System Name                      Device Type
  ------------- -------------------------------- -----------
  0001e6-0421c0 North Sea                        Switch
  000883-07e720 DEFAULT_CONFIG                   Switch
```

MAC addresses of discovered Candidates.

Knowing the available switch numbers (SNs) and Candidate MAC addresses, you can manually assign a Candidate to be a Member of the stack:

Syntax:

stack member *<switch-number>* mac-address *<mac-addr>* [password *<password-str>*]

For example, if the switch in the above output did not have a Manager password and you want to make it a stack Member with an SN of 2, use the following command:

```
switch(config)# stack member 2 mac-address 0060b0-dfla00
```

Using `show stack view` then shows the added Member (see the following example).

Displaying the stack after adding a new member

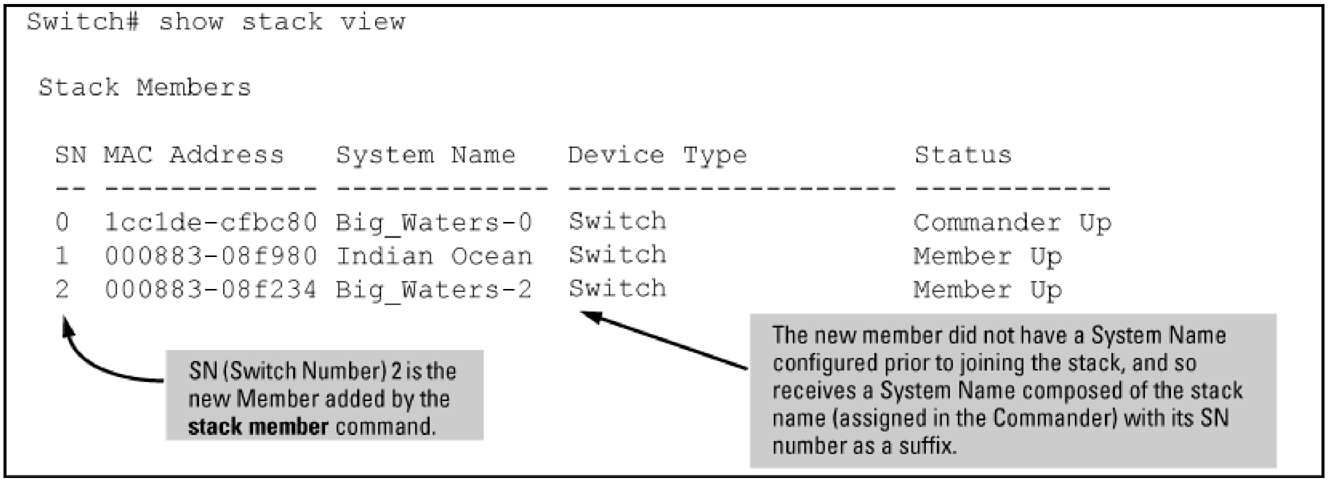

```
Switch# show stack view

 Stack Members

  SN MAC Address    System Name    Device Type           Status
  -- ------------- ------------- -------------------- ------------
  0  1cc1de-cfbc80 Big_Waters-0   Switch                Commander Up
  1  000883-08f980 Indian Ocean   Switch                Member Up
  2  000883-08f234 Big_Waters-2   Switch                Member Up
```

SN (Switch Number) 2 is the new Member added by the **stack member** command.

The new member did not have a System Name configured prior to joining the stack, and so receives a System Name composed of the stack name (assigned in the Commander) with its SN number as a suffix.

Using auto join on a Candidate (CLI)

In the default configuration, a Candidate's Auto Join parameter is set to Yes, meaning that it will automatically join a stack if the stack's Commander detects the Candidate and the Commander's Auto Grab parameter is set to Yes. You can disable Auto Join on a Candidate if you want to prevent automatic joining in this case. There is also the instance where a Candidate's Auto Join is disabled, for example, when a Commander leaves a stack and its members automatically return to Candidate status, or if you manually remove a Member from a stack. In this case, you may want to reset Auto Join to Yes.

Syntax:

```
stack auto-join

no stack auto-join
```

Enables Auto Join on a Candidate.

The `no` version disables Auto Join on a Candidate.

Using a Candidate CLI to push the Candidate into a stack

Use this method if any of the following apply:

- The Candidate's `Auto Join` is set to `Yes` and you do not want to enable `Auto Grab` on the Commander, or the Candidate's `Auto Join` is set to `No`.
- Either you know the MAC address of the Commander for the stack into which you want to insert the Candidate, or the Candidate has a valid IP address and is operating in your network.

Syntax:

```
stack join <mac-addr>
```

`<mac-addr>`

Specifies the MAC address of the Commander in the destination stack.

If the Candidate has an IP address valid for your network use Telnet or a direct serial port connection to access the CLI for the Candidate switch.

Pushing a candidate into a stack

Suppose a Candidate named `North Sea` with `Auto Join` off and a valid IP address of 10.28.227.104 is running on a network. You could Telnet to the Candidate, use `show stack all` to determine the Commander's MAC address, and then push the Candidate into the desired stack.

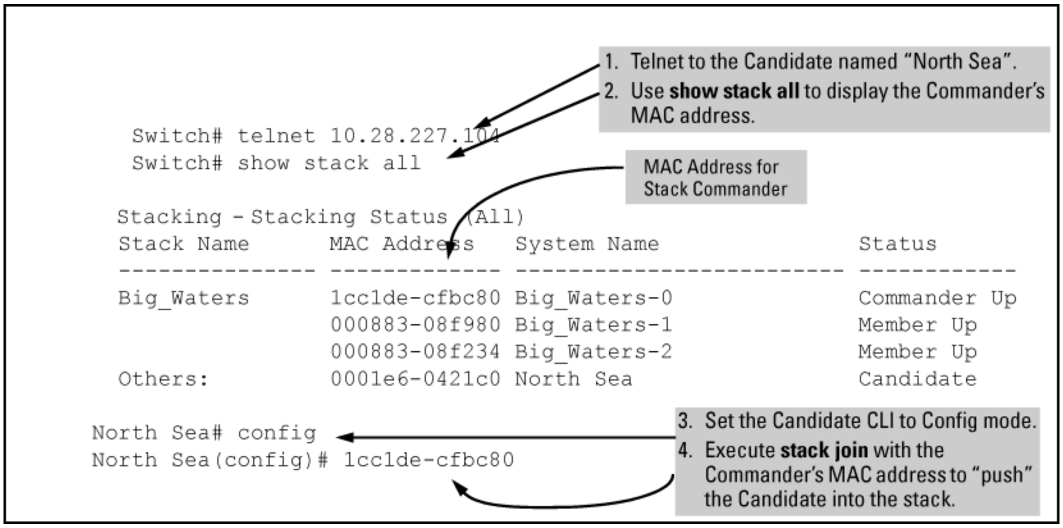

```
Switch# telnet 10.28.227.104
Switch# show stack all

Stacking - Stacking Status (All)
Stack Name      MAC Address   System Name               Status
--------------- ------------- ------------------------- ------------
Big_Waters      1cc1de-cfbc80 Big_Waters-0              Commander Up
                000883-08f980 Big_Waters-1              Member Up
                000883-08f234 Big_Waters-2              Member Up
Others:         0001e6-0421c0 North Sea                 Candidate

North Sea# config
North Sea(config)# 1cc1de-cfbc80
```

1. Telnet to the Candidate named "North Sea".
2. Use **show stack all** to display the Commander's MAC address.

MAC Address for Stack Commander

3. Set the Candidate CLI to Config mode.
4. Execute **stack join** with the Commander's MAC address to "push" the Candidate into the stack.

To verify that the Candidate successfully joined the stack, execute `show stack all` again to view the stacking status.

Using the destination Commander CLI to pull a member from another stack

Syntax:

```
stack member <switch-number> mac-address [password <password-str>]
```

In the destination Commander, finds the MAC address of the Member you want to pull into the destination stack.

Stack output with two stacks in the subnet

Suppose that you create a new Commander with a stack name of `Cold_Waters` and want to move a switch named `Bering Sea` into this new stack:

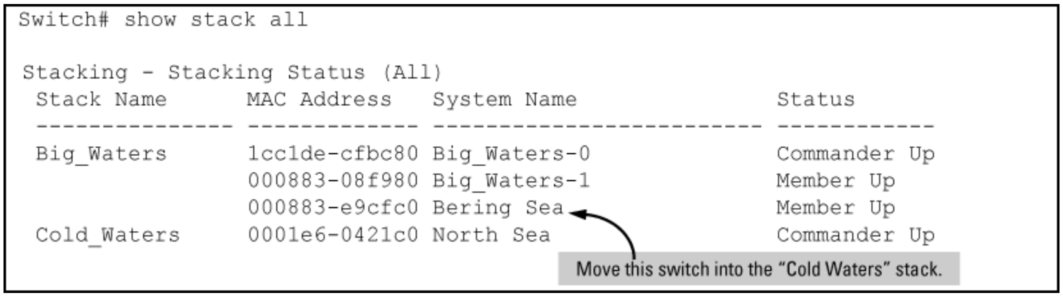

```
Switch# show stack all

Stacking - Stacking Status (All)
 Stack Name      MAC Address   System Name               Status
 --------------- ------------- ------------------------- ------------
 Big_Waters      1cc1de-cfbc80 Big_Waters-0              Commander Up
                 000883-08f980 Big_Waters-1              Member Up
                 000883-e9cfc0 Bering Sea                Member Up
 Cold_Waters     0001e6-0421c0 North Sea                 Commander Up
```

Move this switch into the "Cold Waters" stack.

You would then execute the following command to pull the desired switch into the new stack:

```
switch(config)# stack member 1 mac-address 0060b0-df1a00
```

Where `1` is an unused switch number (`SN`).

Since a password is not set on the Candidate, a password is not needed in this example.

Use `show stack all` again to verify that the move took place.

Using a Member CLI to push the Member into another stack

Use the Member's CLI to push a stack Member into a destination stack if you know the MAC address of the destination Commander.

Syntax:

```
stack join <mac-addr>
```

<mac-addr>

Specifies the MAC address of the Commander in the destination stack.

Converting a Commander to a Member of another stack (CLI)

Syntax:

```
no stack name <stack name>
```

```
stack join <mac-address>
```

Removes the Commander from a stack, eliminates the stack, and returns its Members to the Candidate pool with `Auto Join` disabled.

On using the `no stack` command on a commander, all members will continue to show the commander entry for 125 seconds. After the countdown timer expires on each member for the commander, the members will no longer show the commander entry.

To identify the MAC address of the destination Commander, use the `show stack all` command.

Converting a Commander to a Member

Suppose you have a switch operating as the Commander for a temporary stack named Test. To eliminate the temporary Test stack and convert the switch into a member of an existing stack named `Big_Waters`, execute the following commands in the switch's CLI:

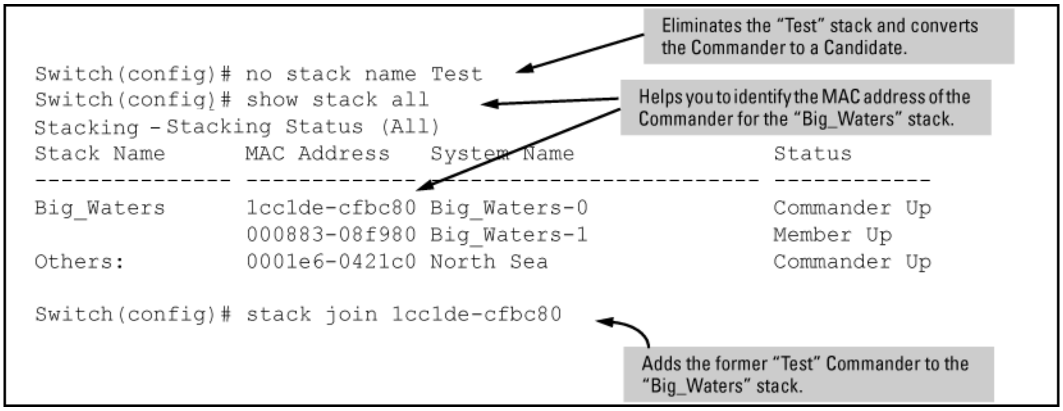

```
Switch(config)# no stack name Test
Switch(config)# show stack all
Stacking - Stacking Status (All)
Stack Name       MAC Address   System Name               Status
---------------- ------------- ------------------------- ------------
Big_Waters       1cc1de-cfbc80 Big_Waters-0              Commander Up
                 000883-08f980 Big_Waters-1              Member Up
Others:          0001e6-0421c0 North Sea                 Commander Up

Switch(config)# stack join 1cc1de-cfbc80
```

Eliminates the "Test" stack and converts the Commander to a Candidate.

Helps you to identify the MAC address of the Commander for the "Big_Waters" stack.

Adds the former "Test" Commander to the "Big_Waters" stack.

Removing a Member from a stack (CLI)

You can remove a Member from a stack using the CLI of either the Commander or the Member.

NOTE:

When you remove a Member from a stack, the Member's `Auto Join` parameter is set to `No`.

Removing a stack Member using the Commander's CLI

This option requires the switch number (SN) and the MAC address of the switch to remove. Because the Commander propagates its Manager password to all stack members, knowing the Manager password is necessary only for gaining access to the Commander. Use `show stack view` to list the stack Members.

Syntax:

```
no stack member <switch-num> mac-address <mac-addr>
```

Removes the specified Member from the stack.

A commander and three switches in a stack

Suppose you want to use the Commander to remove the `North Sea` member from the following stack:

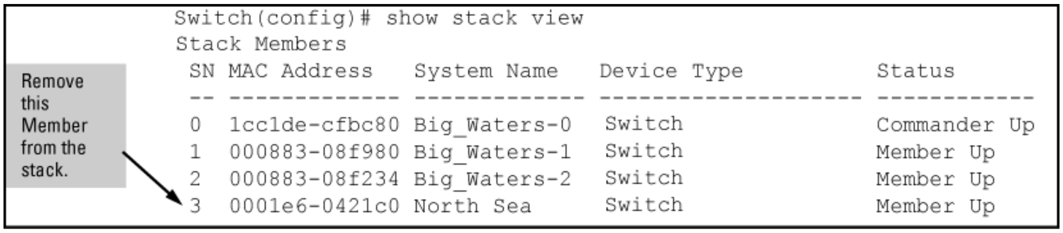

```
Switch(config)# show stack view
Stack Members
 SN MAC Address   System Name   Device Type          Status
 -- ------------- ------------- -------------------- ------------
 0  1cc1de-cfbc80 Big_Waters-0   Switch               Commander Up
 1  000883-08f980 Big_Waters-1   Switch               Member Up
 2  000883-08f234 Big_Waters-2   Switch               Member Up
 3  0001e6-0421c0 North Sea      Switch               Member Up
```

Execute this command to remove the `North Sea` switch from the stack:

```
switch(config)# no stack member 3 mac-address 0030c1-7fc700
```

where:

- 3 is the `North Sea` member's switch number (`SN`)
- `0030c1-7fc700` is the `North Sea` member's MAC address

Removing a stack Member using the Member's CLI

Syntax:

```
no stack join <mac-addr>
```

To use this method you need the Commander's MAC address, which is available using the `show stack` command in the Member's CLI.

Identifying the commander's MAC address from a member switch

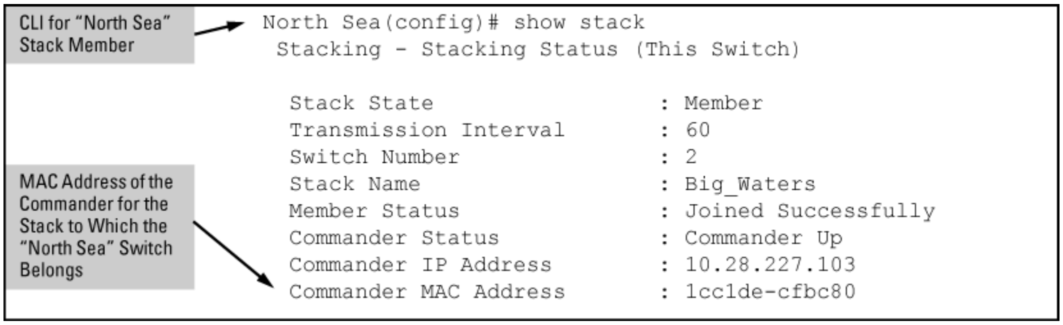

```
North Sea(config)# show stack
 Stacking - Stacking Status (This Switch)

  Stack State                    : Member
  Transmission Interval          : 60
  Switch Number                  : 2
  Stack Name                     : Big_Waters
  Member Status                  : Joined Successfully
  Commander Status               : Commander Up
  Commander IP Address           : 10.28.227.103
  Commander MAC Address          : 1cc1de-cfbc80
```

CLI for "North Sea" Stack Member
MAC Address of the Commander for the Stack to Which the "North Sea" Switch Belongs

Execute the following command in the `North Sea` switch's CLI to remove the switch from the stack:

```
North Sea(config)# no stack join 0030c1-7fec40
```

Accessing Member switches for configuration changes and traffic monitoring (CLI)

After a Candidate becomes a Member, you can use the Telnet command from the Commander to access the Member's CLI or console interface for the same configuration and monitoring you would do through a Telnet or direct-connect access from a terminal.

Syntax:

```
telnet <switch-number>
```

<switch-number>

Specifies an unsigned integer assigned by the Commander to each member (range: `1 - 15`).

To find the switch number for the Member you want to access, execute the `show stack view` command in the Commander's CLI.

A stack showing switch number (SN) assignments

Suppose you want to configure a port trunk on the switch named `North Sea` in the stack named `Big_Waters`. To do so go to the CLI for the `Big_Waters` Commander and execute `show stack view` to find the switch number for the `North Sea` switch:

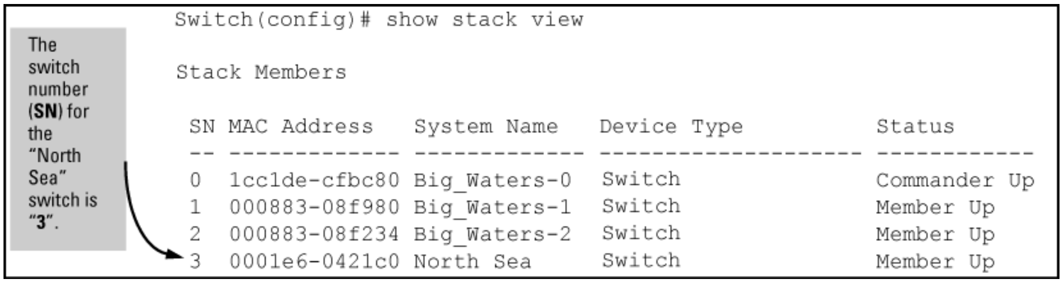

```
Switch(config)# show stack view

Stack Members

 SN MAC Address    System Name    Device Type          Status
 -- -------------- -------------- -------------------- ------------
 0  1cc1de-cfbc80 Big_Waters-0   Switch               Commander Up
 1  000883-08f980 Big_Waters-1   Switch               Member Up
 2  000883-08f234 Big_Waters-2   Switch               Member Up
 3  0001e6-0421c0 North Sea      Switch               Member Up
```

To access the `North Sea` console, execute the following Telnet command:

```
switch(config)# telnet 3
```

You see the CLI prompt for the `North Sea` switch, allowing you to configure or monitor the switch as if you were directly connected to the console.

Disabling or re-enabling stacking (CLI)

In the default configuration stacking is enabled on the switch. You can use the CLI to disable stacking on the switch at any time. Disabling stacking has the following effects:

Disabling a Commander

Eliminates the stack, returns the stack Members to Candidates with `Auto Join` disabled, and changes the Commander to a stand-alone (non-stacking) switch. You must re-enable stacking on the switch before it can become a Candidate, Member, or Commander.

Disabling a Member

Removes the Member from the stack and changes it to a stand-alone (non-stacking) switch. You must re-enable stacking on the switch before it can become a Candidate, Member, or Commander.

Disabling a Candidate

Changes the Candidate to a stand-alone (non-stacking) switch.

Syntax:

```
no stack
```

Disables stacking on the switch.

```
stack
```

Enables stacking on the switch.

Setting the transmission interval (CLI)

All switches in a stack must be set to the same transmission interval to help ensure proper stacking operation. Hewlett Packard Enterprise recommends that you leave this parameter set to the default 60 seconds.

Syntax:

```
stack transmission-interval <seconds>
```

Using the Commander to manage the stack

The Commander normally operates as your stack manager and point of entry into other switches in the stack. This typically includes:

- Adding new stack members
- Moving members between stacks
- Removing members from a stack
- Accessing stack members for individual configuration changes and traffic monitoring

The Commander also imposes its passwords on all stack members and provides SNMP community membership to the stack.

About stack management

Components of Switch stack management

Table 30: *Stacking definitions*

Term	Definition
Stack	Consists of a Commander switch and any Member switches belonging to that Commander's stack.
Commander	A switch that has been manually configured as the controlling device for a stack. When this occurs, the switch's stacking configuration appears as `Commander`.
Candidate	A switch that is ready to join (become a Member of) a stack through either automatic or manual methods. A switch configured as a Candidate is not in a stack.
Member	A switch that has joined a stack and is accessible from the stack Commander.

Figure 41: *A switch moving from Candidate to Member*

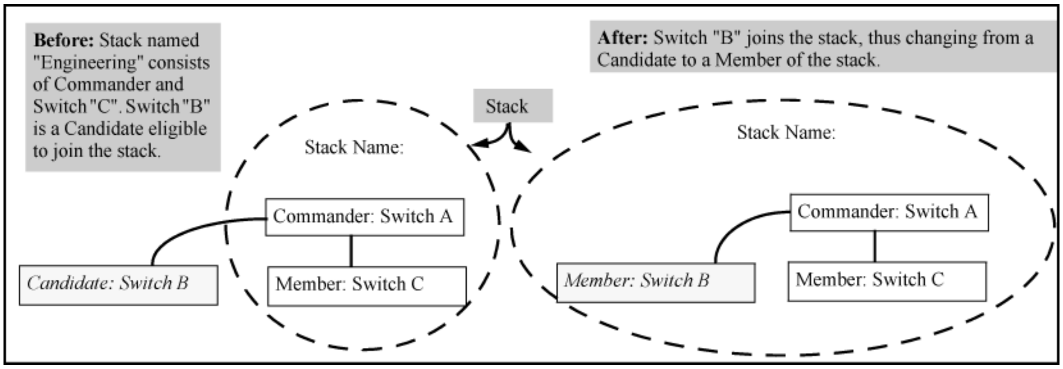
Before: Stack named "Engineering" consists of Commander and Switch "C". Switch "B" is a Candidate eligible to join the stack.
After: Switch "B" joins the stack, thus changing from a Candidate to a Member of the stack.
Stack
Stack Name:
Commander: Switch A
Candidate: Switch B
Member: Switch C
Stack Name:
Commander: Switch A
Member: Switch B
Member: Switch C

General stacking operation

After you configure one switch to operate as the Commander of a stack, additional switches can join the stack by either automatic or manual methods. After a switch becomes a Member, you can work through the Commander switch to further configure the Member switch as necessary for all of the additional software features available in the switch.

The Commander switch serves as the in-band entry point for access to the Member switches. For example, the Commander's IP address becomes the path to all stack Members and the Commander's Manager password controls access to all stack Members.

Figure 42: *Stacking with one commander controlling access to wiring closet switches*

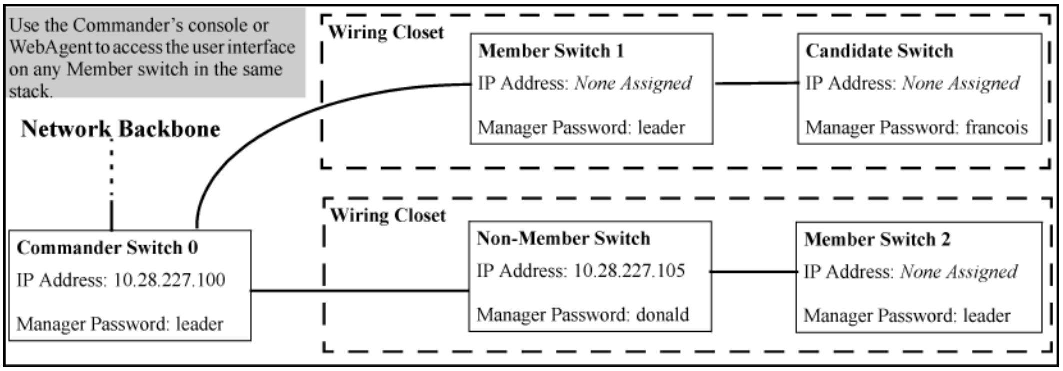
Use the Commander's console or WebAgent to access the user interface on any Member switch in the same stack.
Network Backbone
Commander Switch 0
IP Address: 10.28.227.100
Manager Password: leader
Wiring Closet
Member Switch 1
IP Address: None Assigned
Manager Password: leader
Candidate Switch
IP Address: None Assigned
Manager Password: francois
Wiring Closet
Non-Member Switch
IP Address: 10.28.227.105
Manager Password: donald
Member Switch 2
IP Address: None Assigned
Manager Password: leader

Interface options

You can configure stacking through the switch's menu interface, CLI, or the WebAgent For information on how to use the WebAgent to configure stacking, see the online help by clicking on the **?** in the WebAgent screen.

Operating rules for stacking

General rules

- Stacking is an optional feature (enabled in the default configuration) and can easily be disabled. Stacking has no effect on the normal operation of the switch in your network.
- A stack requires one Commander switch. (Only one Commander allowed per stack.)
- All switches in a particular stack must be in the same IP subnet (broadcast domain). A stack cannot cross a router.
- A stack accepts up to 16 switches (numbered 0-15), including the Commander (always numbered 0).
- The stacking feature supports up to 100 switches in the same IP subnet (broadcast domain), however, a switch can belong to only one stack. In the event that the 100 switch limit is exceeded, it may take multiple attempts to add or move a member to any given stack. Once a member is added to a stack, it is not "forgotten" by the Commander.
- The `stack status (all)` command will display up to 100 devices. Devices that are not members of a given stack may periodically drop out of the list.

- If multiple VLANs are configured, stacking uses only the primary VLAN on any switch. In the factory-default configuration, the DEFAULT_VLAN is the primary VLAN. (See **Stacking operation with multiple VLANs configured** on page 266 and **The primary VLAN** on page 53.)
- Stacking allows intermediate devices that do not support stacking. This enables you to include switches that are distant from the Commander.

Figure 43: *A non-stacking device used in a stacking environment*

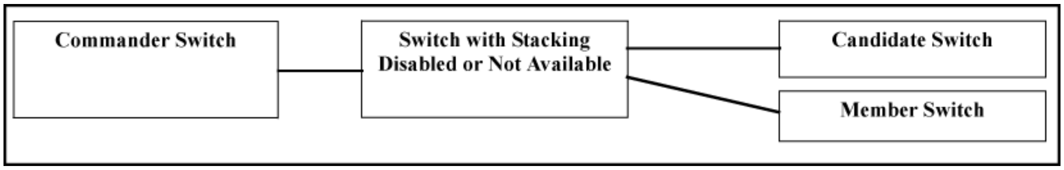

Specific rules for commander, candidate, and member switch

Table 31: *Specific rules for commander, candidate, and member switch*

	IP Addressing and Stack Name	Number Allowed Per Stack	Passwords	SNMP Communities
Commander	`IP Addr`: Requires an assigned IP address and mask for access via the network. `Stack Name`: Required	Only one Commander switch is allowed per stack.	The Commander's Manager and Operator passwords are assigned to any switch becoming a Member of the stack. If you change the Commander's passwords, the Commander propagates the new passwords to all stack Members.	Standard SNMP community operation. The Commander also operates as an SNMP proxy to Members for all SNMP communities configured in the Commander.
Candidate	`IP Addr`: Optional. Configuring an IP address allows access via Telnet or WebAgent while the switch is not a stack member. In the factory default configuration the switch automatically acquires an IP address if your network includes DHCP service. `Stack Name`: N/A	n/a	Passwords optional. If the Candidate becomes a stack Member, it assumes the Commander's Manager and Operator passwords. If a candidate has a password, it cannot be automatically added to a stack. In this case, if you want the Candidate in a stack, you must manually add it to the stack.	Uses standard SNMP community operation if the Candidate has its own IP addressing.
Member	`IP Addr`: Optional. Configuring an IP address allows access via Telnet or WebAgent without going through the Commander switch. This is useful, for example, if the stack Commander	Up to 15 Members per stack.	When the switch joins the stack, it automatically assumes the Commander's Manager and Operator passwords and discards any passwords it may	Belongs to the same SNMP communities as the Commander (which serves as an SNMP proxy to the Member for communities to which the Commander

	IP Addressing and Stack Name	Number Allowed Per Stack	Passwords	SNMP Communities
	fails and you need to convert a Member switch to operate as a replacement Commander. Stack Name: N/A		have had while a Candidate. Note: If a Member leaves a stack for any reason, it retains the passwords assigned to the stack Commander at the time of departure from the stack.	belongs). To join other communities that exclude the Commander, the Member must have its own IP address. Loss of stack membership means loss of membership in any community that is configured only in the Commander.

NOTE: In the default stack configuration, the Candidate Auto Join parameter is enabled, but the Commander Auto Grab parameter is disabled. This prevents Candidates from automatically joining a stack prematurely or joining the wrong stack (if more than one stack Commander is configured in a subnet or broadcast domain). If you plan to install more than one stack in a subnet, Hewlett Packard Enterprise recommends that you leave Auto Grab disabled on all Commander switches and manually add Members to their stacks. Similarly, if you plan to install a stack in a subnet (broadcast domain) where stacking-capable switches are not intended for stack membership, you should set the Stack State parameter (in the Stack Configuration screen) to Disabled on those particular switches.

Stacking operation with multiple VLANs configured

Stacking uses the primary VLAN in a switch. In the factory-default configuration, the DEFAULT_VLAN is the primary VLAN. However, you can designate any VLAN configured in the switch as the primary VLAN. (See **The primary VLAN** on page 53.)

When using stacking in a multiple-VLAN environment, the following criteria applies:

- Stacking uses only the primary VLAN on each switch in a stack.
- The primary VLAN can be tagged or untagged as needed in the stacking path from switch to switch.
- The same VLAN ID (VID) must be assigned to the primary VLAN in each stacked switch.

Status messages

Stacking screens and listings display these status messages:

Message	Condition	Action or Remedy
Candidate Auto-join	Indicates a switch configured with Stack State set to `Candidate`, `Auto Join` set to `Yes` (the default) and no Manager password.	None required
Candidate	Candidate cannot automatically join the stack because one or both of the following conditions apply: • Candidate has `Auto Join` set to `No`. • Candidate has a Manager password.	Manually add the candidate to the stack.
Commander Down	Member has lost connectivity to its Commander.	Check connectivity between the Commander and the Member.
Commander Up	The Member has stacking connectivity with the Commander.	None required.
Mismatch	This may be a temporary condition while a Candidate is trying to join a stack. If the Candidate does not join, then stack configuration is inconsistent.	Initially, wait for an update. If the condition persists, re-configure the Commander or the Member.
Member Down	A Member has become detached from the stack. A possible cause is an interruption to the link between the Member and the Commander.	Check the connectivity between the Commander and the Member.
Member Up	The Commander has stacking connectivity to the Member.	None required.
Rejected	The Candidate has failed to be added to the stack.	The candidate may have a password. In this case, manually add the candidate. Otherwise, the stack may already be full. A stack can hold up to 15 Members (plus the Commander).

SNMP community operation in a stack

Community Membership

In the default stacking configuration, when a Candidate joins a stack, it automatically becomes a Member of any SNMP community to which the Commander belongs, even though any community names configured in the Commander are not propagated to the Member's SNMP Communities listing. However, if a Member has its own (optional) IP addressing, it can belong to SNMP communities to which other switches in the stack, including the Commander, do not belong. For example:

Figure 44: *An SNMP community operation with stacking*

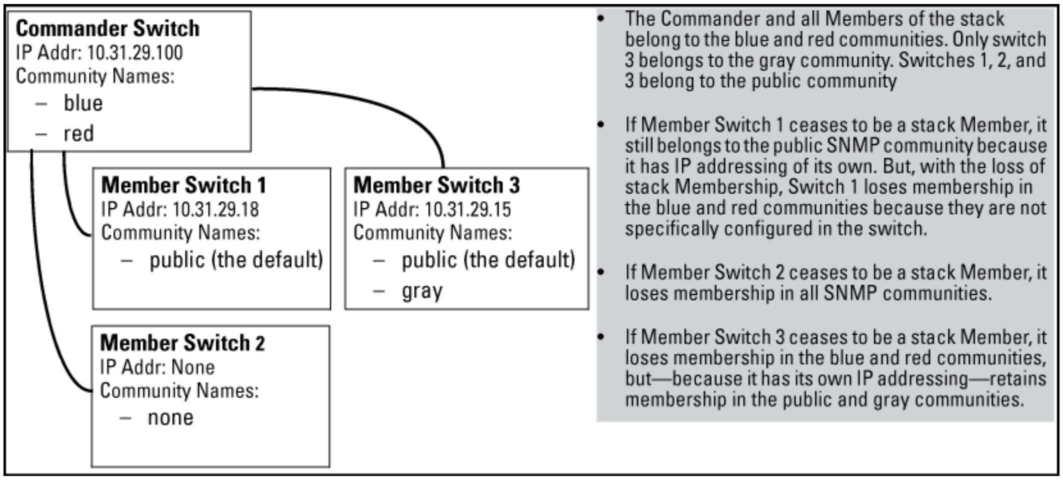

- The Commander and all Members of the stack belong to the blue and red communities. Only switch 3 belongs to the gray community. Switches 1, 2, and 3 belong to the public community
- If Member Switch 1 ceases to be a stack Member, it still belongs to the public SNMP community because it has IP addressing of its own. But, with the loss of stack Membership, Switch 1 loses membership in the blue and red communities because they are not specifically configured in the switch.
- If Member Switch 2 ceases to be a stack Member, it loses membership in all SNMP communities.
- If Member Switch 3 ceases to be a stack Member, it loses membership in the blue and red communities, but—because it has its own IP addressing—retains membership in the public and gray communities.

SNMP management station access to members via the Commander

To use a management station for SNMP Get or Set access through the Commander's IP address to a Member, you must append `@sw <switch number>` to the community name. For example, in the figure in **Community Membership**, you would use the following command in your management station to access Switch 1's MIB using the blue community:

```
snmpget MIB variable 10.31.29.100 blue@sw1
```

Because the gray community is only on switch 3, you could not use the Commander IP address for gray community access from the management station. Instead, you would access switch 3 directly using the switch's own IP address. For example:

```
snmpget MIB variable 10.31.29.15 gray
```

In the figure in **Community Membership**, you cannot use the public community through the Commander to access any of the Member switches. For example, you can use the public community to access the MIB in switches 1 and 3 by using their unique IP addresses. However, you must use the red or blue community to access the MIB for switch 2.

```
snmpget MIB variable 10.31.29.100 blue@sw2
```

Introduction to BYOD-redirect

The BYOD (bring-your-own-device) feature lets you design, manage, and control a BYOD network when you configure the BYOD-redirect feature on your switches.

Where BYOD-redirect is enabled on a switch, the device's client credentials are sent to the BYOD server for registration. The BYOD server stores the registration information for each client's device (such as the device MAC-address), which gives that client's device access to the network.

The BYOD solution includes:

- secure user authentication
- centralized authentication process
- authorization and accounting
- unified monitoring and network management services
- ease-of-use self-registration (on-boarding) process

BYOD solution

The following figure illustrates a BYOD solution that includes the following:

- Access point and wireless controller: manages wireless SSIDs.
- BYOD (IMC) server: manages BYOD policy and centralized user management.

- switches: redirects user registration traffic to IMC and grants access to ports.
- BYOD Redirect feature

Figure 45: *BYOD solution*

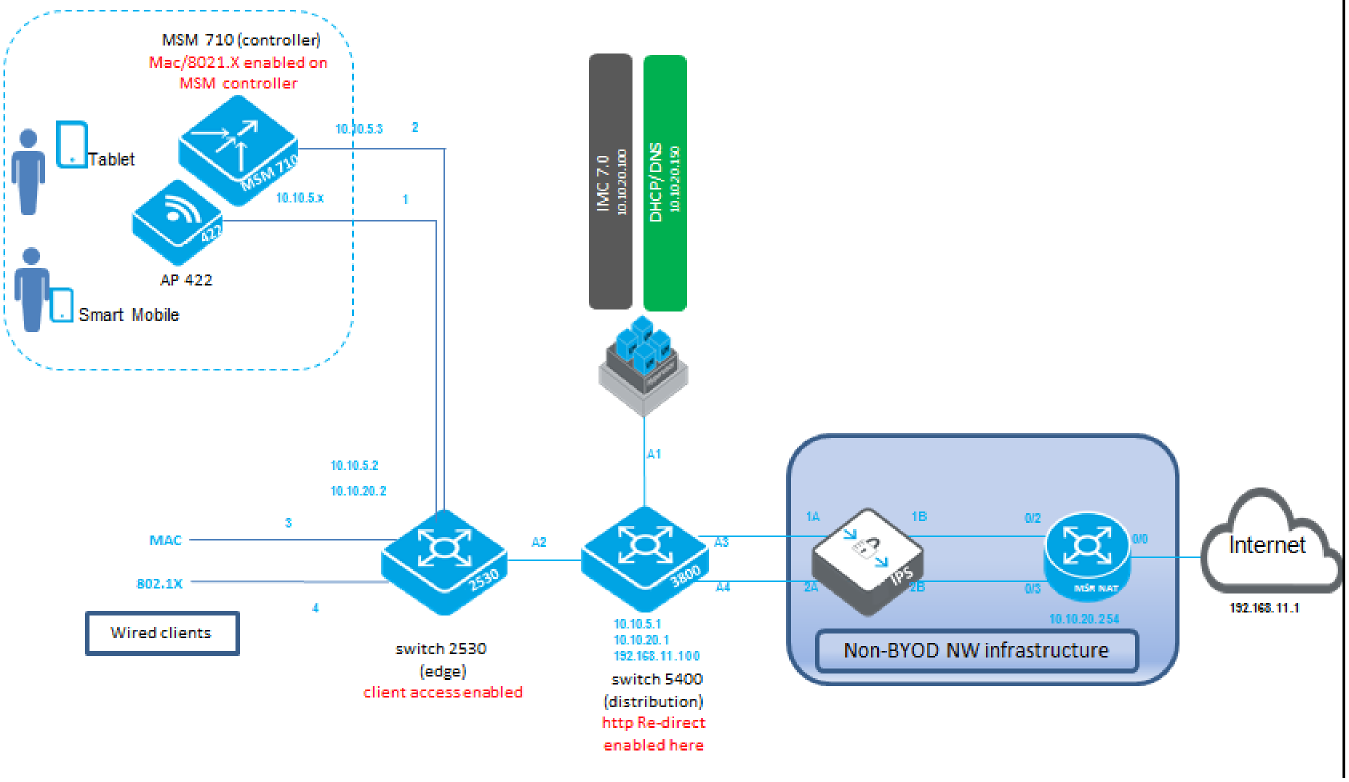

BYOD features

When BYOD-redirect is enabled on a VLAN, the BYOD feature intercepts HTTP traffic and blocks all other traffic for which free rules are not enabled. Most BYOD-redirect implementation is platform independent, except installing free rules to mitigate risks.

Communication between clients and the IMC server is tunneled by the edge switch:

1. A client request is read by the HTTP task.
2. The HTTP task always redirects, after embedding client IP addresses, a URL trying to access the redirected URL.

3. The redirect response includes URL parameters: **user ip address** and **url user is trying to access.**
4. The client receives a redirect response from the switch and makes an HTTP request to redirect the URL.

Figure 46: *The BYOD-redirect function*

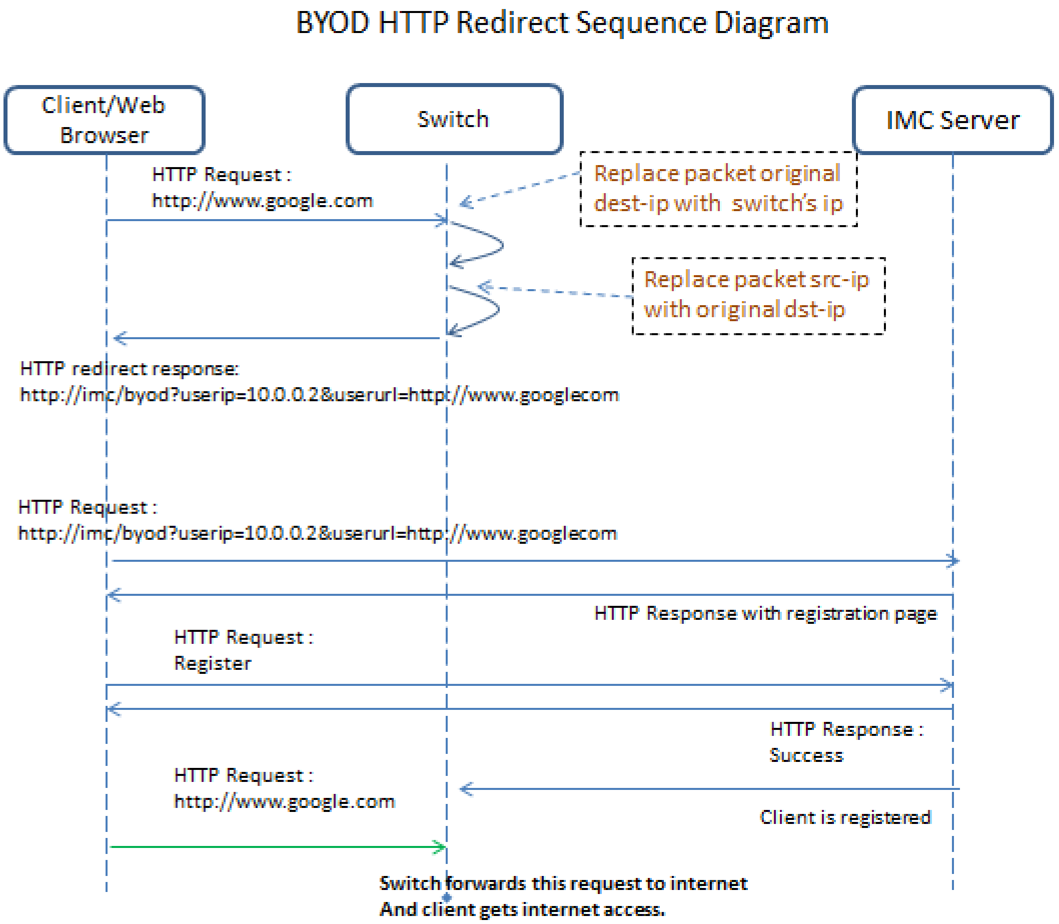
BYOD HTTP Redirect Sequence Diagram
Client/Web Browser
Switch
IMC Server
HTTP Request : http://www.google.com
Replace packet original dest-ip with switch's ip
Replace packet src-ip with original dst-ip
HTTP redirect response: http://imc/byod?userip=10.0.0.2&userurl=http://www.googlecom
HTTP Request : http://imc/byod?userip=10.0.0.2&userurl=http://www.googlecom
HTTP Response with registration page
HTTP Request : Register
HTTP Response : Success
HTTP Request : http://www.google.com
Client is registered
Switch forwards this request to internet
And client gets internet access.

Interoperability with other switch features

The following rules can help avoid conflicts when BYOD-redirect has been deployed on a switch with other features:

1. **MAFR and BYOD-redirect are mutually exclusive** – MAFR (MAC Authentication Failure Redirect) and BYOD-redirect solve similar problems.
2. **DNS sentinel and BYOD-redirect** – When a DNS sentinel is enabled, the switch tunnels packets to the controller. Packets are re-injected to the switch only if the controller classifies DNS packets as permitted. When BYOD-redirect is enabled, the user should configure an ACL rule to pass through DNS packets to the switch. If SDN controller policy classifies a DNS packet originating from a client as drop, then BYOD-redirect does not work.
3. **IP sentinel and BYOD-redirect** – When IP sentinel is enabled for the IP flows configured by the SDN controller, the switch tunnels the IP packets to the controller. The IP packets are reinjected to the switch only if the controller classifies the IP traffic as not malicious. If the SDN controller policy classifies the client's IP traffic as malicious, then BYOD-redirect fails.

4. **OpenFlow and BYOD-redirect** – If an OpenFlow instance is enabled on a VLAN, then all traffic is given to the OpenFlow packet processing task. BYOD-redirect requires intercepting IP (HTTP) packets. If BYOD-redirect interoperates with OpenFlow, traffic should be copied to both OpenFlow and BYOD-redirect; otherwise, the switch cannot enable BYOD-redirect and OpenFlow on the same VLAN.
5. **Other TCAM rules** – If any other user has configured TCAM rules that override TCAM entries installed for BYOD-redirect, BYOD redirect does not work.

Interoperability with other vendors

Because BYOD policy integrates several logical components including MSM, UAM, and RADIUS, the redirected URL in the BYOD-redirect feature on a switch must include the `byod-server-url` and `user-ip` information to work with the IMC server.

Restrictions

BYOD-redirect has the following restrictions:

1. BYOD-redirect is a per-VLAN configuration; up to three VLANs can be enabled with BYOD-redirect.
2. BYOD-redirect supports up to three redirection servers configured on a switch. When a redirection server URL is configured, the BYOD module maintains separate data structures to store the redirected URL on the VLAN where BYOD-redirect is enabled. BYOD-redirect statistics are maintained for each server.

Configuring BYOD

Creating a BYOD server

Configure a portal redirect web-server.

Syntax

```
no portal web-server [web-server-name] url [url-string]
```

`portal`: Configure the BYOD redirect feature.

`web-server`: Configure portal redirect web-server.

`web-server-name`: Specify the BYOD web-server name in ASCII.

`url`: Configure the URL of the BYOD server.

`url-string` : A URL redirecting the client to the BYOD server must be in ASCII.

Associating a BYOD server

Associate a BYOD server with a specific VLAN to redirect clients to the assigned URL page.

Syntax

```
no vlan [vlan-id] <portal web-server [web-server-name]>
```

`vlan`: Add, delete, edit VLAN configuration, or enter a VLAN context.

`vlan-id`: VLAN identifier or VLAN name.

`portal`: Configure the BYOD redirect feature on a VLAN.

`web-server`: Specify the BYOD web-server.

web-server-name: BYOD web-server name in ASCII.

Creating a BYOD ACL rule

Configure a BYOD-free rule.

Syntax

```
no portal free-rule [rule-number] vlan [VLAN-ID] destination <<ip-address> | mask
<mask-length> | any tcp <des-tcp-port> | udp <des-udp-port> | source <ip-address>
| mask <mask-length> | any tcp <src-tcp-port> |udp <src-udp-port>>
```

Term	Meaning
portal	Configure the BYOD redirect feature.
free-rule	Configure a BYOD-free rule.
rule-number	Free rule number as an INTEGER**<1-6>**.
vlan	Free rule source VLAN ID.
VLAN-ID	VLAN identifier or VLAN name.
destination	Free rule destination.
ip-address	IP address
mask	Mask
mask-length	Mask length.
tcp	TCP protocol
udp	UDP Protocol
des-udp-port	tcp port destination
source	Free rule source.
<src/des-tcp/udp-port>	TCP or UDP port number, as an integer**<1-65534>**.
any	Free rule source any.
ip	Free rule source IP.
IP	Free rule destination IP.
any	Free rule source or destination any.

Implementing BYOD-redirect configuration

BYOD enables employees to register and access corporate resources with personally-owned devices. Though BYOD provides flexibility to employees, it can bring challenges to IT departments. BYOD-redirect is designed to help manage and control personal devices and policies at the enterprise network level.

Before implementing BYOD-redirect ensure that:

- BYOD-redirect is configured on a VLAN.
- BYOD-redirect is supported on up to three VLANs.
- BYOD-redirect is supported with Mac and 802.1X authentications.
- BYOD-redirect works with IMC 7.0 UAM module.
- The client URL and DHCP IP are included in the Redirect URL to the IMC.

NOTE:

Until the registration process has been completed, a client device cannot access the internet or the enterprise network. Any traffic from this unauthorized device is redirected to the BYOD server.

Implementing BYOD-redirect configuration examples

The following examples show how to implement BYOD-redirect for both wired and wireless solutions.

BYOD configuration on a distribution switch

To facilitate the BYOD-redirect function, complete the following tasks on the distribution switch:

1. Configure DNS and make FQDN solution successful: `ip dns server-address priority 1` <**DNS-server-IP**>.

   NOTE: The argument to the URL can be an FQDN or IP address. If you use the IP address as an argument, this step is not necessary.

2. Configure BYOD web-server URL: **portal web-server "byod" url http://imc.com:8080/byod**.
3. Enable BYOD-redirect on a VLAN: **vlan 101 portal web-server "byod."**
4. Configure BYOD-redirect free-rules on the on-boarding VLAN 101 to permit client traffic transit through DNS and DHCP servers using the following commands.To permit DNS traffic to/from a DNS server to a client through on-boarding VLAN:

   a. `portal free-rule 1 vlan 101 source any udp 0 destination any udp 53`

   b. `portal free-rule 2 vlan 101 source any udp 53 destination any udp 0`

   To permit DHCP traffic to/from DHCP server to client through on-boarding VLAN:

   a. `portal free-rule 3 vlan 101 source any udp 68 destination any udp 67`

   b. `portal free-rule 4 vlan 101 source any udp 67 destination any udp 68`

5. Register the device in IMC on the on-boarding VLAN. When registration is successful, client traffic is placed into different VLAN (guest/corporate) configurations.

Client authentication configuration on edge switch

Enable MAC authentication on edge switch port 1-2 using the following commands:

- `# enable mac authentication on ports 1-2`
- `aaa port-access mac-based 1-2`
- `# configure number of client limits on port 1 and port2`
- `aaa port-access mac-based 1 addr-limit 32`
- `aaa port-access mac-based 2 addr-limit 32`
- `radius-server host <radius ip> dyn-authorization`
- `radius-server host <radius ip> time-window 0`

Table 32: *Wired and wireless components configured in a network topology*

Access Type	Edge Switch	Distribution Switch	Configuration ProcedureNote
Wired Access	Edge switch (for example 2530)	5400 switch	**1.** Register the edge switch in HPE IMC. **2.** Create the configuration on the edge switch. **3.** Create the configuration on 5400 switch.
Wireless Access			**1.** Make the HPE MSM controller reachable by IMC. **2.** Ensure that access points (HPE 422) are managed by the MSM controller. **3.** Configure MAC or 802.1X authentication on the MSM controller. **4.** Create the configuration on the 5400 switch.

Figure 47: *Wired and wireless components configured in a network topology*

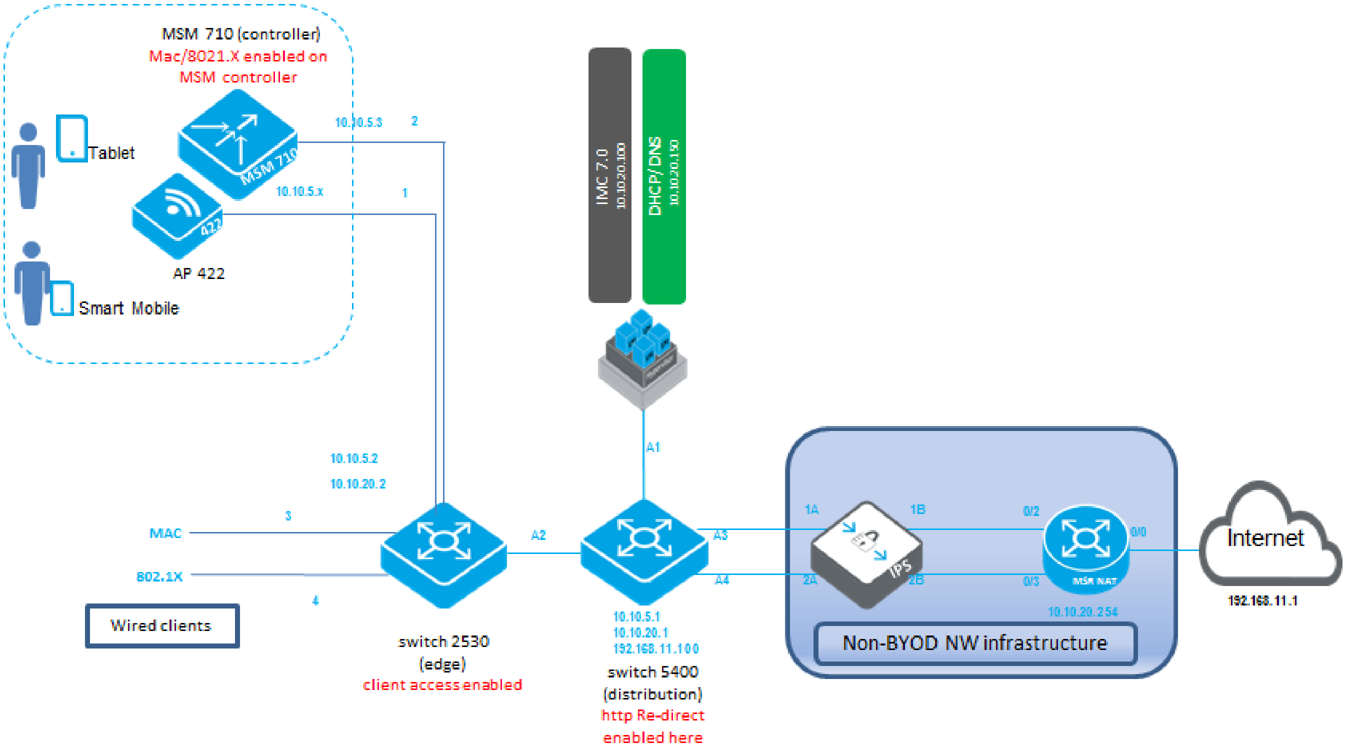

Table 33: *Wired clients solution*

Access Type	Edge Switch	Distribution Switch	Configuration Procedure
Wired Access	Edge switch (for example 2530)	Switch 3810	**1.** Register the edge switch and distribution switch in IMC. **2.** Ensure that both the edge and distribution switch can reach the DHCP and DNS server. **3.** Create the configuration on the edge switch. **4.** Create the configuration on the distribution switch.

Figure 48: *Wired clients solution*

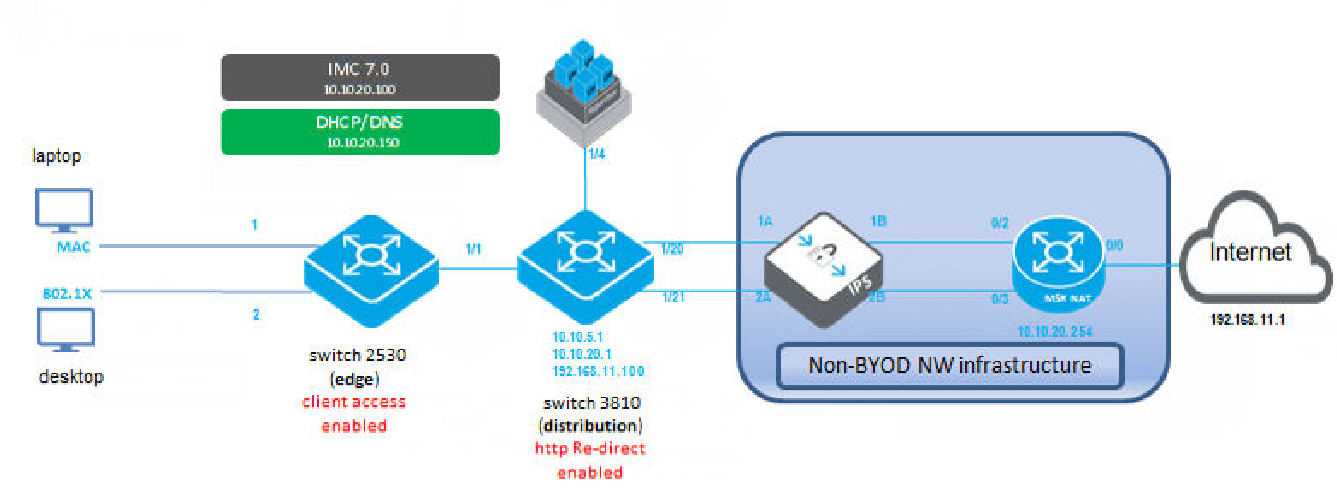

Table 34: *Configuration and access for wired clients on an edge switch*

Access Type	Edge Switch	Distribution Switch	Configuration Procedure
Wired Access	Edge switch	N/A	**1.** Register the edge switch in IMC. **2.** Ensure that the edge switch is reachable by the DHCP and DNS server. **3.** Create the configuration on the edge switch. **4.** Create the configuration on the edge switch.

Figure 49: *Configuration and access for wired clients on an edge switch*

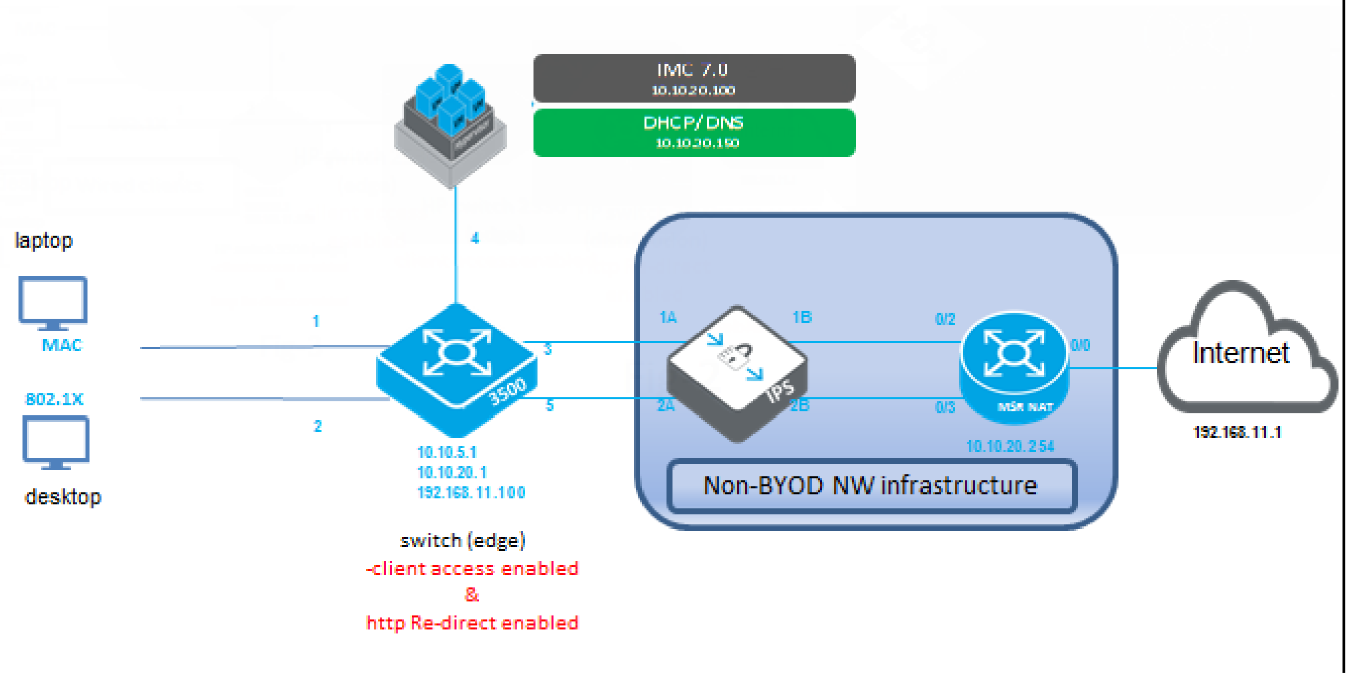

Show commands

Show portal server

Display all BYOD servers and their attributes or specify a BYOD web-server-name to display its details.

Syntax

```
show portal web-server [web-server-name]
```

Term	Meaning
portal	Display BYOD server details..
web-server	Specify the BYOD web-server.
web-server name	Enter BYOD web-server name in ASCII.

Sample output

```
Portal Server:
1)imc:
Resolved IP       : 15.146.197.224
VPN Instance      : n/a
URL               : http://15.146.197.224:80/byod
VLAN              : 101
DNS Cache Status : 20 seconds
```

Show portal redirect statistics

Show redirect statistics of a BYOD.

Syntax

```
show portal redirect statistics
```

Term	Meaning
portal	Display BYOD server details.
redirect	Display redirect statistics
statistics	Display the statistics.

Sample output

```
show portal redirect statistics
Status and Counters - Portal Redirect Information
Total Opens              : 0
Resets Connections       : 0
Current Opens            : 0
Packets Received         : 14997
Packets Sent             : 12013
HTTP Packets Sent        : 3002
Current Connection States :
SYN_RECVD      : 0
ESTABLISHED    : 0
```

Show portal free rule

Display all BYOD free rules and their attributes; the user can specify a BYOD rule to display its free rule.

Syntax

```
show portal free-rule [free-rule-number]
```

Term	Meaning
portal	Display BYOD server details.
free-rule	Display BYOD-free rule.
free-rule-number	Free rule number as an integer <**0-50**>.

Sample output

```
Rule-Number   : 2
Vlan          : 0
Source:
Protocol   : UDP
Port       : 12345
IP         : 0.0.0.0
Mask       : 0.0.0.0
MAC        : n/a
Interface : n/a
Destination:
Protocol   : UDP
Port       : 123
IP         : 0.0.0.0
Mask       : 0.0.0.0
```

Associating with the BYOD server on a specified VLAN

Associate a BYOD server with a specific VLAN to redirect clients to the assigned URL page.

Syntax

```
no vlan <VLAN-ID > [portal web-server < web-server-name>]
```

Term	Meaning
portal	Configure the BYOD redirect feature on the VLAN.
web-server	Specify the BYOD web-server.
ASCII-STR	BYOD web server name.
vlan	Add, delete, edit VLAN configuration, or enter a VLAN context.
VLAN-ID	Enter a VLAN identifier or a VLAN name.

Overview of smart link

Smart link is a switch feature that provides effective, simple, and fast-converging link redundancy in network topology with dual uplink between different layers of the network. It requires an active (master) and a backup (slave) link. The active link carries the uplink traffic. Upon failure of the active link, a switchover is triggered and the traffic is directed to the backup link.

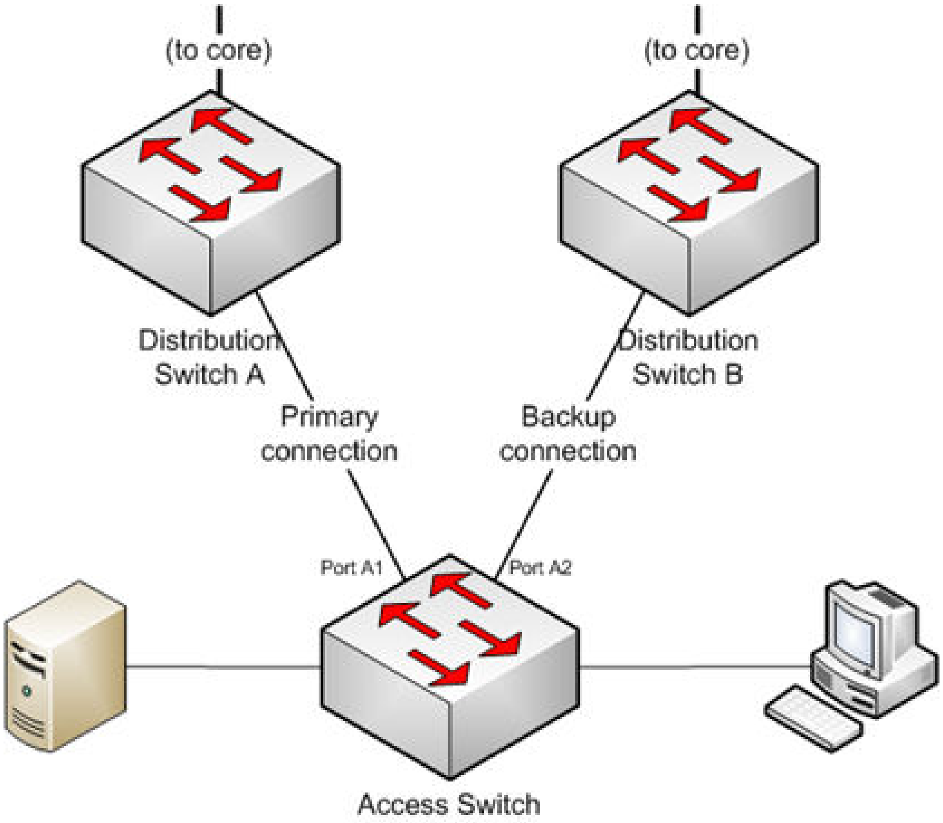

- In the figure above, ports A1 and A2 are configured as part of a smart link group. The connection from the access switch to Distribution Switch A is the master, and the connection from the access switch to Distribution Switch B is the slave.
- Only the master interface forwards traffic for a group of vlans (called a protected vlan group).
- The other interface is in standby mode for this protected group. If port A1 goes down, port A2 starts forwarding traffic for this protected vlan group.
- If port A1 comes back up, it goes to standby mode and does not forward traffic. Port A2 continues forwarding traffic. This is the case if preemption-mode is configured as “role”. If preemption-mode is not configured as “role”, when the master (A1) comes back up, it becomes Active (forwarding) after the configured ‘preemption-delay’.
- Since a smart link group has it’s configuration information readily available for which port should be forwarding for the protected vlan group in the case of the active link failure, failover is much quicker than with STP.

Smart link configuration commands

Create a smart link group

Syntax

```
no smart link group group-id
```

Create a smart link group. When command is entered without any parameter, it enters into smart link group context.

`master` ***`port`***

Assign master port.

`slave` ***`port`***

Assign slave port.

`protected-vlans` ***`vid-list`***

Assign protected VLANs.

`send-control-vlan` ***`vid`***

Assign the VLAN to send flush packets.

`preemption-mode` ***`off`***

Specify the preemption mode. (Default is off.)

`preemption-delay` ***`10-max`***

Set the delay until when standby preempts active. (Default is 1 second.)

`trap` ***`enable`*** **`|`** ***`disable`***

Enable sending trap for this group.

Disable the trap for this group

NOTE: The maximum number of Smartlink Groups supported is 24.

Configure VLANs

Syntax

```
no smart link recv-control-vlan <vid-list>
```

Configures VLANs to receive flush messages. This is interface level command. Command must be executed for both master and slave port.

Enable debug

Syntax

```
no debug smart link group group-id | all flush-packets
```

Enable debug messages for a smart link group.

Configuration example

The following example illustrates smart link configuration with VLAN load-balancing:

```
      vlans 1-10 mapped to smart link group 1
smart link group 1 master a1 slave a2
smart link group 1 protected-vlans 1-10
smart link group 1 send-control-vlan 1
smart link group 1 preemption-mode role
smart link group 1 preemption-delay 10
```

```
      vlans 11-20 mapped to smart link group 2
smart link group 2 master a2 slave a1
smart link group 2 protected-vlans11-20
smart link group 2 send-control-vlan 10
smart link group 2 preemption-mode role
smart link group 2 preemption-delay 15
```

Show smart link group

Show the smart link group information. Detailed output is displayed if group is specified, otherwise only basic information is displayed for all groups.

Syntax

```
show smart link group group 1-24 | all | flush-statistics | recv-control-vlans
```

Show smart link group information.

flush-statistics

Show information about the received flush messages.

group

Show information for groups.

recv-control-vlans

Show receive control VLANs information.

show smart link group all

```
Switch# show smart link group all

Smart link Group Information:

       Master Slave   Active   Standby   Protected   Send        Preemption   Preemption
 Grp.  Port   Port    Port      Port       VLANs     Ctrl VLAN   Mode          delay
----- ------ ------- --------- -------- ---------- ----------- ----------- ----------
1       A1      A2      A1        A2        1-5        10         Role          10
2       A4      A5      A4        A5        6           7         Off           10

Switch#show smart link group 1
Smartlink Group 1 Information:
Protected VLANs       : 1-5
Preemption Mode [Off] : Role
Send Control VLAN     : 10
Preemption Delay      : 10
Trap        [Disable] : Enable
```

```
Ports  Role   State    Flush Count Last Flush Time
------ ------ -------  ----------- --------------------
1      Master Active   5           16:37:20 2013/06/17
2      Slave  Standby  5           16:37:20 2013/06/17
```

Show smart link flush-statistics

Show statistics of received flush packets.

Syntax

```
show smart link flush-statistics
```

The show command helps to display information about the received flush messages.

show smart link flush-statistics

```
Switch# show smart link flush-statistics

Last Flush Packet Detail:
Flush Packets Received                      : 2
Last Flush Packet Received On Interface     : 23
Last Flush Packet Received On               : 00:11:07 1990/01/01
Device Id Of Last Flush Packet Received     : c8cbb8-ddc0c0
Control VLAN Of Last Flush Packet Received : 1
```

Show receive control

Syntax

```
show smart link recv-control-vlans
```

Show receive control VLANs configured on per port basis.

show smart link recv-control-vlan

```
Switch# show smart link recv-control-vlan

 Receive Control VLAN Information:
 Port         VLANs
 -----     --------------
 A1              1-3
 B1              4
```

Show tech smart link

Syntax

```
show tech all|custom|buffers|instrumentation|mesh|route|route stale|stat|vrrp|smart link|transparentmode smart link
```

Display output of a predefined command sequence used by technical support.

show tech smart link

```
show tech smart link
```

```
Smartlink Group 1 Information:
Protected VLANs        : 1-5
Send Control VLAN      : 10
Preemption Mode [Off] : Role
Preemption Delay       : 10
Trap          [Disable] : Enable

Ports  Role   State    Flush Count Last Flush Time
------ ------ -------- ----------- --------------------
1      Master Active   5           16:37:20 2013/06/17
2      Slave  Standby  5           16:37:20 2013/06/17

show smart link flush
Last Flush Packet Detail:
Flush Packets Received                     : 2
Last Flush Packet Received On Interface    : 23
Last Flush Packet Received On              : 00:11:07 1990/01/01
Device Id Of Last Flush Packet Received    : c8cbb8-ddc0c0
Control VLAN Of Last Flush Packet Received : 1
```

Clear command

Clear group and flush statistics

Syntax

```
clear smart link flush-statistics group group-id | all
```

Event Log

Event	Message
Whenever a standby port transits to active port.	Port A1 is now active on smart link group 10

Chapter 13 Websites

Networking Websites

Aruba Support Portal

asp.arubanetworks.com

Aruba Software and Documentation

asp.arubanetworks.com/downloads

Hewlett Packard Enterprise Networking Software

www.hpe.com/networking/software

Hewlett Packard Enterprise Networking website

www.hpe.com/info/networking

Hewlett Packard Enterprise My Networking website

www.hpe.com/networking/support

Hewlett Packard Enterprise My Networking Portal

www.hpe.com/networking/mynetworking

Hewlett Packard Enterprise Networking Warranty

www.hpe.com/networking/warranty

General websites

Hewlett Packard Enterprise Information Library

www.hpe.com/info/EIL

For additional websites, see **Support and other resources**.

Accessing Hewlett Packard Enterprise Support

- For live assistance, go to the Contact Hewlett Packard Enterprise Worldwide website:

  http://www.hpe.com/info/assistance

- To access documentation and support services, go to the Hewlett Packard Enterprise Support Center website:

  http://www.hpe.com/support/hpesc

Information to collect

- Technical support registration number (if applicable)
- Product name, model or version, and serial number
- Operating system name and version
- Firmware version
- Error messages
- Product-specific reports and logs
- Add-on products or components
- Third-party products or components

Accessing updates

- Some software products provide a mechanism for accessing software updates through the product interface. Review your product documentation to identify the recommended software update method.

- To download product updates:

  Hewlett Packard Enterprise Support Center

  www.hpe.com/support/hpesc

  Hewlett Packard Enterprise Support Center: Software downloads

  www.hpe.com/support/downloads

  Software Depot

  www.hpe.com/support/softwaredepot

- To subscribe to eNewsletters and alerts:

  www.hpe.com/support/e-updates

- To view and update your entitlements, and to link your contracts and warranties with your profile, go to the Hewlett Packard Enterprise Support Center **More Information on Access to Support Materials** page:

  www.hpe.com/support/AccessToSupportMaterials

IMPORTANT: Access to some updates might require product entitlement when accessed through the Hewlett Packard Enterprise Support Center. You must have an HPE Passport set up with relevant entitlements.

Customer self repair

Hewlett Packard Enterprise customer self repair (CSR) programs allow you to repair your product. If a CSR part needs to be replaced, it will be shipped directly to you so that you can install it at your convenience. Some parts do not qualify for CSR. Your Hewlett Packard Enterprise authorized service provider will determine whether a repair can be accomplished by CSR.

For more information about CSR, contact your local service provider or go to the CSR website:

http://www.hpe.com/support/selfrepair

Remote support

Remote support is available with supported devices as part of your warranty or contractual support agreement. It provides intelligent event diagnosis, and automatic, secure submission of hardware event notifications to Hewlett Packard Enterprise, which will initiate a fast and accurate resolution based on your product's service level. Hewlett Packard Enterprise strongly recommends that you register your device for remote support.

If your product includes additional remote support details, use search to locate that information.

Remote support and Proactive Care information

HPE Get Connected

www.hpe.com/services/getconnected

HPE Proactive Care services

www.hpe.com/services/proactivecare

HPE Datacenter Care services

www.hpe.com/services/datacentercare

HPE Proactive Care service: Supported products list

www.hpe.com/services/proactivecaresupportedproducts

HPE Proactive Care advanced service: Supported products list

www.hpe.com/services/proactivecareadvancedsupportedproducts

Proactive Care customer information

Proactive Care central

www.hpe.com/services/proactivecarecentral

Proactive Care service activation

www.hpe.com/services/proactivecarecentralgetstarted

Warranty information

To view the warranty information for your product, see the links provided below:

HPE ProLiant and IA-32 Servers and Options

www.hpe.com/support/ProLiantServers-Warranties

HPE Enterprise and Cloudline Servers

www.hpe.com/support/EnterpriseServers-Warranties

HPE Storage Products

www.hpe.com/support/Storage-Warranties

HPE Networking Products

www.hpe.com/support/Networking-Warranties

Regulatory information

To view the regulatory information for your product, view the *Safety and Compliance Information for Server, Storage, Power, Networking, and Rack Products*, available at the Hewlett Packard Enterprise Support Center:

www.hpe.com/support/Safety-Compliance-EnterpriseProducts

Additional regulatory information

Hewlett Packard Enterprise is committed to providing our customers with information about the chemical substances in our products as needed to comply with legal requirements such as REACH (Regulation EC No 1907/2006 of the European Parliament and the Council). A chemical information report for this product can be found at:

www.hpe.com/info/reach

For Hewlett Packard Enterprise product environmental and safety information and compliance data, including RoHS and REACH, see:

www.hpe.com/info/ecodata

For Hewlett Packard Enterprise environmental information, including company programs, product recycling, and energy efficiency, see:

www.hpe.com/info/environment

Documentation feedback

Hewlett Packard Enterprise is committed to providing documentation that meets your needs. To help us improve the documentation, send any errors, suggestions, or comments to Documentation Feedback (**docsfeedback@hpe.com**). When submitting your feedback, include the document title, part number, edition, and publication date located on the front cover of the document. For online help content, include the product name, product version, help edition, and publication date located on the legal notices page.